

A symplectic Gysin sequence

TIMOTHY PERUTZ

We use the theory of pseudo-holomorphic quilts to establish a counterpart, in symplectic Floer homology, to the Gysin sequence for the homology of a sphere-bundle. In a motivating class of examples, this ‘symplectic Gysin sequence’ is precisely analogous to an exact sequence describing the behaviour of Seiberg–Witten monopole Floer homology for 3-manifolds under connected sum.

1 Introduction

1.1 Spherically-fibred coisotropic submanifolds

This paper is about the following situation. We are given

- compact symplectic manifolds (M, ω_M) of dimension $2n$ and (N, ω_N) of dimension $2(n - k)$, where $0 \leq k \leq n$;
- a smooth fibre bundle $\rho: V \rightarrow N$ whose typical fibre F is a k -dimensional manifold admitting a Morse function with precisely two critical points (so F is homeomorphic to S^k);
- a smooth codimension- k embedding $i: V \hookrightarrow M$, such that

$$(1) \quad i^* \omega_M = \rho^* \omega_N.$$

We permit M to have non-empty boundary provided that $i(V) \subset \text{int}(M)$ and that the boundary is of contact type with outward-pointing Liouville field. However, N must be a closed manifold.

More economically, we can suppose that the data given are (M, ω_M) and its submanifold V , which must be a *spherically fibred coisotropic submanifold* (a notion which we shall define momentarily). The remaining data, (N, ω_N) and ρ , can then be reconstructed. Recall that a coisotropic submanifold $V \subset M$ is one such that the annihilator $\text{Ann}(TV)$ of TV with respect to ω_M is contained in TV . As each point $x \in V$, ω_M then pairs $\text{Ann}(T_x V)$ non-degenerately with $T_x M / T_x V$, hence $\text{Ann}(TV)$ has constant rank

k . Moreover, V is foliated by isotropic leaves [18, Lemma 5.33] with k -dimensional tangent spaces $\text{Ann}(T_x V)$, and there is a quotient map ρ from V to the leaf-space N . Say that V is *fibred coisotropic*¹ if the isotropic foliation is fibrating: that is, each leaf F has an open neighbourhood U for which there is diffeomorphism $U \cong F \times D^k$ taking the isotropic foliation to the product foliation. When this happens, N has induced smooth and symplectic structures, and ρ becomes a smooth submersion satisfying (1).

Given $(M, \omega_M; N, \omega_N; V, \rho)$ satisfying (1), there is a natural symplectomorphism of (N, ω_N) with the leaf-space (or reduced symplectic manifold) of V , intertwining $\rho: V \rightarrow N$ with the quotient map from V to its leaf-space.

A fibred coisotropic submanifold is *spherically fibred* when the quotient map to its leaf-space is a bundle whose fibres are homeomorphic to the sphere S^k .

One reason to regard fibred coisotropics as a natural class of submanifolds is that, unlike general coisotropic submanifolds, they have a straightforward deformation theory [25]. Modulo Hamiltonian diffeomorphisms, the infinitesimal deformations are parametrised by $H^0(N; \mathcal{H}^1)$ where \mathcal{H}^1 is the local system on N associated with (V, ρ) with fibres $H^1(\rho^{-1}(x); \mathbb{R})$.

A simple property of fibred coisotropic submanifolds is the following:

Proposition 1.1 *The graph $\widehat{V} := \text{gr}(\rho)$ is Lagrangian in $(M, -\omega_M) \times (N, \omega_N)$.*

Proof If (u_1, v_1) and (u_2, v_2) lie in $T_{(x, \rho(x))} V \subset T_x M \times T_{\rho(x)} N$ then $v_1 = \rho_* u_1$ and $v_2 = \rho_* u_2$, hence $\omega_M(u_1, u_2) = (\rho^* \omega_N)(u_1, u_2) = \omega_N(v_1, v_2)$. Thus \widehat{V} is isotropic, and by a dimension-count it is Lagrangian. Alternatively, even in infinite dimensions, we can argue that if $(u, \rho_* u') \in \text{Ann}(T_{(x, \rho(x))} \widehat{V})$ then for all $w \in T_x V$ one has $\omega_M(u, w) = \omega_N(\rho_* u', \rho_* w) = (\rho^* \omega_N)(u', w) = \omega_M(u', w)$; hence $u - u' \in \text{Ann}(T_x V) = \ker(D\rho)$, so $\rho_* u' = \rho_* u$ and \widehat{V} is coisotropic. \square

Notation 1.2 *We usually indicate that we wish to consider the symplectic form $-\omega_M$ on M , rather than ω_M , by writing M_- . We distinguish fibred coisotropic submanifolds from their associated Lagrangian correspondences by adding a hat in the latter case, as in the previous proposition.*

The subject of this paper is the Lagrangian Floer homology

$$(2) \quad \text{HF}(\widehat{V}, (\text{id}_{M_-} \times \mu)\widehat{V}),$$

taken inside $M_- \times N$, where $\mu \in \text{Aut}(N, \omega_N)$. Our aim is to establish a relationship between this and the fixed-point Floer homology $\text{HF}(\mu)$. This takes the form of a long exact sequence, the ‘symplectic Gysin sequence’ of the title.

¹Or ‘integral coisotropic’, in W. Ruan’s terminology [25].

1.2 Statement of results

To ensure that our Floer homology groups and the maps between them are well-defined we impose monotonicity hypotheses on the symplectic manifolds M and N and the Lagrangian correspondence \widehat{V} . A connected symplectic manifold (M, ω) is *monotone* if the the symplectic area homomorphism $\pi_2(M) \rightarrow \mathbb{R}$, defined by $f \mapsto \int_{S^2} f^* \omega$, is a non-negative multiple of the the Chern class homomorphism $f \mapsto \int_{S^2} f^* c_1(TM, \omega)$. A compact connected Lagrangian submanifold $L \subset M$ is monotone if the symplectic area $\pi_2(M, L) \rightarrow \mathbb{R}$ is a non-negative multiple of the Maslov index m_L . A symplectic manifold containing a monotone Lagrangian is necessarily monotone, since the image of $c_1(TM, \omega)$ in $\text{Hom}(\pi_2(M, L), \mathbb{R})$ is $2m_L$. The *minimal Maslov index* m_L^{\min} of a Lagrangian is the non-negative generator of the additive subgroup $m_L(\pi_2(M, L)) \subset \mathbb{Z}$.

Hypothesis 1.3 (MON) \widehat{V} is monotone in $M_- \times N$ with $m_{\widehat{V}}^{\min} \geq 2$.

The hypothesis MON is sufficient to guarantee that $\text{HF}(\widehat{V}, (\text{id}_M \times \mu)\widehat{V})$ is well-defined, with mod 2 Novikov coefficients (when $m_{\widehat{V}}^{\min} = 2$ this follows from the argument of [21, Addendum]) but not strong enough to validate the symplectic Gysin sequence. We supplement it with the following:

Hypothesis 1.4 (MAS) $m_{\widehat{V}}^{\min} \geq k + 2$.

Notation 1.5 The universal Novikov ring Λ_R of the commutative ring R is the ring of formal series $\sum_{r \in \mathbb{R}} a(r)t^r$, with $a(r) \in R$, such that $\text{supp}(a) \cap (-\infty, s]$ is finite for each $s \in \mathbb{R}$.

Theorem 1 (GYSIN SEQUENCE) Assume that MON and MAS hold. For each $\mu \in \text{Aut}(N, \omega_N)$ there are then canonical maps forming an exact triangle

$$\begin{array}{ccc}
 \text{HF}(\mu) & \xrightarrow{e(V) \frown \cdot} & \text{HF}(\mu) \\
 & \swarrow \rho_* \quad \nwarrow \rho^! & \\
 & \text{HF}(\widehat{V}, (\text{id}_{M_-} \times \mu)\widehat{V}) &
 \end{array}$$

Here $e(V) \in H^{k+1}(N; R)$ is the Euler class of the S^k -bundle $\rho: V \rightarrow N$, and the horizontal arrow is its action on fixed-point Floer homology by quantum cap product. All three terms in the triangle are taken with coefficients in the universal Novikov ring Λ_R of a commutative ring R of characteristic 2.

Remarks (i) All three maps in the sequence have a definite degree, but since Floer-theoretic gradings are relative, only the horizontal arrow has a straightforward *numerical* degree. This degree is $-k-1$. If one declares that ρ_* has degree 0 (modulo $m_{\widehat{V}}^{\min}$) then $\rho^!$ has degree k . Note that $(-k-1) + 0 + k = -1$.

(ii) The hypotheses are sharp in that there is no odd k for which MAS could be weakened to $m_{\widehat{V}}^{\min} \geq k+1$ (see Example 1.13). In Theorem 6.2, we analyse the borderline case, $m_{\widehat{V}}^{\min} = k+1$. We show that the sequence is valid when the map $e(V) \frown$ is replaced by $x \mapsto e(V) \frown x + \nu t^a x$, where a is the area of the discs of Maslov index $k+1$, t the formal Novikov variable, and ν is the mod 2 count of index $k+1$ pseudo-holomorphic discs attached to \widehat{V} , with two marked boundary points mapping respectively to a point and to a *global angular chain*.

Addendum 1.6 *The symplectic Gysin sequence has the following additional properties.*

- The maps $\rho^!$ and ρ_* intertwine the natural actions of the quantum cohomology $QH^*(N; \Lambda_{\mathbb{Z}/2})$.
- One can lift the characteristic two assumption and take, for example, $R = \mathbb{Z}$, providing that (V, ρ) is an R -orientable bundle and that either $k = 1$, or there is a vector bundle $E \rightarrow M$ such that $w_1(E) = 0$ and $w_2(E)|_V = w_2(\ker D\rho)$.

If we are in a situation where $\mathrm{HF}_*(\widehat{V}, \widehat{V}) \cong H_*(\widehat{V}; \Lambda_R)$, we obtain from the case $\mu = \mathrm{id}_N$ an exact triangle which resembles the classical Gysin sequence:

$$(3) \quad \begin{array}{ccc} H_*(N; \Lambda_R) & \xrightarrow{e(V) \frown_q \cdot} & H_{*-k-1}(N; \Lambda_R) \\ & \swarrow \scriptstyle [-1] \quad \searrow & \\ & H_{*-1}(V; \Lambda_R) & \end{array}$$

The three maps in the classical Gysin sequence are $e(V) \frown \cdot$, $\rho^!$ and ρ_* (where, of course, these have their classical meanings). We have stated that the horizontal map in our sequence is the quantum cap product $e(V) \frown_q \cdot$, a deformation of the classical cap product $e(V) \frown \cdot$.

Addendum 1.7 *The map $\rho^!$ is a deformation of the arrow \swarrow from diagram 3. That is, \swarrow is the lowest order part of the map $\rho^!$; the remaining part has strictly positive Novikov weight.*

The proofs of the addenda are routine. However, a further refinement remains unproven. One has $\mathrm{HF}_*(\mu) = \mathrm{HF}^*(\mu^{-1})$ and $\mathrm{HF}_*(\widehat{V}, (\mathrm{id} \times \mu)\widehat{V}) = \mathrm{HF}^*(\widehat{V}, \mathrm{id} \times \mu^{-1}\widehat{V})$ ('Poincaré duality'). Conjecture: *under Poincaré duality, $\rho^!$ (for μ) corresponds to the adjoint of ρ_* (for μ^{-1}).* We will see that $\rho^!$ is defined geometrically, and is associated a picture (Figure 3), while ρ_* is defined algebraically. The conjecture is that ρ_* should be associated with a similar picture, with the orientation reversed.

There is another hypothesis under which one has an exact Gysin sequence: that of *strong negativity*.

Definition 1.8 A symplectic manifold (X^{2n}, η) is *strongly negative* if every class $\beta \in \pi_2(X)$ with $\langle \beta, [\eta] \rangle > 0$ satisfies $\langle \beta, c_1(TX) \rangle \leq -n - 2$. A compact Lagrangian L is strongly negative if every $\delta \in \pi_2(X, L)$ with $\langle \delta, [\eta] \rangle > 0$ satisfies $m_L(\delta) \leq -n - 2$.

Theorem 1.9 *The terms in the Gysin sequence are well-defined when M and N are strongly negative and $\widehat{V} \subset M_- \times N$ is a strongly negative Lagrangian. Moreover, Theorem 1 and its two addenda continue to hold in this situation.*

1.3 Examples and corollaries

The basic example of a spherically-fibred coisotropic submanifold is the following.

Example 1.10 Let $S = S^{2n-1}(r)$ be a sphere of radius r in (standard symplectic) \mathbb{C}^n . It is coisotropic, and the isotropic leaves are the orbits under the scaling S^1 -action; thus the reduced manifold is \mathbb{CP}^{n-1} and the reduction map $S \rightarrow \mathbb{CP}^{n-1}$ is the Hopf fibration. The symplectic form on \mathbb{CP}^{n-1} is $r\omega_{\mathrm{FS}}$. Thus there are Lagrangian spheres in $M_- \times \mathbb{CP}^{n-1}$, for any symplectic manifold M^{2n} (we have to take r small enough to find a Darboux chart of radius r).

Two mechanisms for producing fibred coisotropic submanifolds are *reduction* and *degeneration*.

Symplectic reduction. If the compact group G acts by Hamiltonian diffeomorphisms of M via a moment map μ , and $\mu^{-1}(s)$ is a regular level set on which G acts freely, then $\mu^{-1}(s)$ is a fibred coisotropic with G as fibre. In particular, when G is $\mathrm{U}(1)$ or $\mathrm{SU}(2)$, $\mu^{-1}(s)$ is spherically fibred.

Degeneration. Consider a symplectic manifold (E, Ω) and a proper map $\pi: E \rightarrow S$ to a Riemann surface S . Assume that the critical set $N := \mathrm{crit}(\pi)$ is smooth and contained in some fibre E_s ; that there is an almost complex structure J in a neighbourhood of

N making π holomorphic (so that N is an almost complex submanifold) such that the Hessian form $D^2\pi$ on its normal bundle, $\nu = \nu_{N/E}$, is non-degenerate as a complex quadratic form on each fibre ν_x . Let $\gamma: [0, 1] \rightarrow S$ be a smooth embedding with $\gamma(1) = s$, and put $M = E_{\gamma(0)}$. Parallel transport over $\gamma|_{[0,t]}$ defines a symplectomorphism $\rho_t: M = E_{\gamma(0)} \rightarrow E_{\gamma(t)}$, for any $t \in [0, 1)$. The *vanishing cycle*

$$V_\gamma := \{x \in M : \lim_{t \rightarrow 1} \rho_t(x) \in N\}$$

is then a spherically-fibred coisotropic submanifold. The fibres are copies of S^k , where $2k + 2$ is the real codimension of N in E , and the reduction map is the limiting parallel transport $\rho = \lim_{t \rightarrow 1} \rho_t: V_\gamma \rightarrow N$. The bundle ρ comes with a natural reduction of structure to $\mathrm{SO}(k + 1)$. Conversely, any spherically-fibred coisotropic submanifold with structure group reduced to $\mathrm{SO}(k + 1)$ arises in this way. For further details, see [24, Section 2].

Example 1.11 Let (M, ω) be a compact symplectic manifold, and fix $p \in M$. Let $\widetilde{M \times \mathbb{C}}$ be a symplectic blow-up of $M \times \mathbb{C}$ near $(p, 0)$. Then the blow-down map $\widetilde{M \times \mathbb{C}} \rightarrow M \times \mathbb{C}$ followed by second projection defines a degeneration $\widetilde{M \times \mathbb{C}} \rightarrow \mathbb{C}$ of the type just described (with $k = 1$). The smooth fibres are symplectomorphic to M ; the zero-fibre is a normal crossing variety $\widetilde{M} \cup_{\mathbb{CP}^{n-1}} \mathbb{CP}^n$, so its singular locus N is a copy of \mathbb{CP}^{n-1} . The normal bundle ν is $\mathcal{O}(1) \oplus \mathcal{O}(-1)$, and the Hessian form is the natural pairing of $\mathcal{O}(1)$ with $\mathcal{O}(-1)$. The vanishing cycle (for the path $[-1, 0]$) is just the sphere $S^{2n-1} \subset M$ of Example 1.10. Its radius in a Darboux chart is the weight of the blow-up, i.e., the area of a projective line \mathbb{CP}^1 in the exceptional divisor.

Remark (i) The ‘complex Morse–Bott’-type degenerations described above occur routinely in algebraic geometry. The case $k = 1$ is especially prevalent, because this is the situation which occurs when the singular fibre E_s has a smooth normal crossing singularity. At the end of the paper we shall take up the examples studied in [24], in which M is a symmetric product of a Riemann surface.

(ii) The symmetric product examples just mentioned arise in a context related to 3- and 4-dimensional topology. The Floer homology groups we shall consider, as well as the maps between them, appear as part of a conjectural symplectic model for Seiberg–Witten gauge theory (see [24]). We shall explain that, for this class of examples, our exact sequence is precisely analogous to a connected sum formula in Seiberg–Witten monopole Floer homology.

Taking $\mu = \mathrm{id}_N$, the symplectic Gysin sequence gives a method for computing Lagrangian Floer *self*-homology. The standard ‘PSS’ isomorphism $\mathrm{HF}(\mathrm{id}_N) \cong H_*(N; \Lambda_R)$

respects the actions on the two sides by the (small) quantum cohomology ring of N [19, Section 12.1]. One can compare Theorem 1 with a result of Oh about Lagrangian Floer self-homology (which in the monotone case may be defined over any commutative ring A of characteristic 2 without the need for Novikov rings):

Theorem 1.12 (Oh [22]) *If L is a monotone Lagrangian with $m_L^{\min} \geq \dim(L) + 2$, one has $\mathrm{HF}(L, L; A) \cong H_*(L; A)$.*

If N is a point, so that $\dim(V) = n = k$, hypothesis MAS coincides with Oh's $m_V^{\min} \geq \dim(V) + 2$. On the other hand, if N is positive-dimensional, MAS is a weaker assumption than Oh's. Our result implies that if the quantum cap product on N is undeformed, or if $e(V) = 0$, one has $\mathrm{HF}_*(\widehat{V}, \widehat{V}; \Lambda_R) \cong H_*(V; \Lambda_R)$.

Example 1.13 Consider the spherical Lagrangian $\widehat{S} = S^{2n-1} \subset (\mathbb{C}^n)_- \times \mathbb{CP}^{n-1}$ from Example 1.10. One has $m_{\widehat{S}}^{\min} = 2n = \dim(\widehat{S}) + 1$. But \widehat{S} is Hamiltonian displaceable, by (translation) $\times \mathrm{id}_{\mathbb{CP}^{n-1}}$, so $\mathrm{HF}_*(\widehat{S}, \widehat{S}) = 0$. If we take $M = D^{2n}(0; R)_- \times \mathbb{CP}^{n-1}$ with $R \gg 0$, and $N = \{\mathrm{pt.}\}$, MAS is, reassuringly, not fulfilled. One sees from this example that hypothesis MAS is sharp when $k = 2n - 1 = \frac{1}{2} \dim(M)$.

If in the same example, with $n \geq 2$, we take M to be the $2n$ -disc $D^{2n}(0; R)$ and $N = \mathbb{CP}^{n-1}$, then $k + 2 = 3 < 2n$, so the Gysin sequence *does* apply. The Euler class e is the hyperplane class h , and quantum cap product by h is invertible; indeed, the quantum cap operator $H = h \frown_q \cdot$ on $H_*(\mathbb{CP}^{n-1}; \Lambda_R)$ has as inverse $t^{-1}H^{n-1}$, where $t \in \Lambda_R$ is the formal Novikov parameter (see, e.g., [18]). Thus Theorem 1 correctly predicts that $\mathrm{HF}_*(\widehat{S}, \widehat{S}) = 0$ (or, reversing one's viewpoint, it shows that in $QH^*(\mathbb{CP}^{n-1})$ one has a relation $h^{*n} = f(t)$ for some $f \neq 0$).

The last example points to a corollary of Theorem 1.

Corollary 1.14 *When MON and MAS hold, and (V, ρ) is R -orientable, \widehat{V} is Hamiltonian displaceable in $M_- \times N$ only if $e(V)$ acts invertibly on the quantum homology $QH_*(N; \Lambda_R)$.*

Indeed, this is the criterion for the vanishing of $\mathrm{HF}_*(\widehat{V}, \widehat{V})$. The unusual aspect of this corollary is that it invokes only *closed* holomorphic curves.

It has recently been shown by Biran and Cornea [3] that, for a large class of monotone Lagrangians L , there is a dichotomy: either the Floer self-homology (over $\Lambda_{\mathbb{Z}/2}$) vanishes, or else it is isomorphic to $H_*(L; \Lambda_{\mathbb{Z}/2})$. Their hypotheses are that $m_L^{\min} \geq 2$ and that the cohomology ring $H^*(L; \mathbb{Z}/2)$ is generated in degrees $\leq m_L^{\min} - 1$. This, together with our vanishing criterion, has the following corollary.

Theorem 1.15 *Suppose that the hypotheses of Theorem 1 hold, and that $H^*(\widehat{V}; \mathbb{Z}/2)$ is generated in degrees $\leq m_{\widehat{V}} - 1$. Then either $e(V)$ acts invertibly on $QH_*(N; \Lambda_{\mathbb{Z}/2})$, or $HF_*(\widehat{V}, \widehat{V}; \Lambda_{\mathbb{Z}/2}) \cong H_*(\widehat{V}; \Lambda_{\mathbb{Z}/2})$.*

1.4 Remarks on the proof

One approach to Theorem 1, along the lines of a standard proof of its classical counterpart, would be to study the spectral sequence associated with a natural filtration on $CF_*(\widehat{V}, (\text{id}_{M_-} \times \mu)\widehat{V})$. A generator for this complex projects to a fixed point of μ , and the filtration is by the Floer-theoretic (relative) degree of this fixed point, considered as a generator of $CF_*(\mu)$. The task is to identify the E^2 -page and the higher differentials.

We shall proceed by a different (more geometric?) route, exhibiting the sequence as the exact triangle associated with a mapping cone. Under hypothesis MON, we shall construct a chain map

$$(4) \quad C\rho: CF_*(\mu) \rightarrow CF_*(\widehat{V}, (\text{id}_{M_-} \times \mu)\widehat{V}).$$

When MAS also holds the composite, quantum cap $e(V) \frown \cdot$ followed by $C\rho$, is nullhomotopic; by defining a nullhomotopy we shall construct a chain map

$$(5) \quad C\Phi: \text{cone} \left(CF_*(\mu) \xrightarrow{e(V) \frown \cdot} CF_{*-k-1}(\mu) \right) \rightarrow CF_*^{M_- \times N}(\widehat{V}, (\text{id}_{M_-} \times \mu)\widehat{V}).$$

Theorem 2 *The map $C\Phi$ induces an isomorphism Φ on homology.*

Theorem 1 is an immediate consequence, since the homology of the mapping cone sits inside a canonical exact triangle.

The classical Gysin sequence can be proved by considering the chain map

$$(6) \quad C\Psi: \text{cone} \left(S_*(N) \xrightarrow{\tilde{e} \frown \cdot} S_*(N)[-k-1] \right) \rightarrow S_*(V)[-k-1],$$

$$(\sigma, \sigma') \mapsto \psi \frown \rho^! \sigma + \rho^! \sigma'.$$

Here $S_*(N)$ and $S_*(V)$ are singular chain complexes, \tilde{e} a cochain representative for $e(V)$, and $\psi \in S^k(V)$ a cochain satisfying $d\psi = \rho^* \tilde{e}$ and restricting to each fibre as an orientation cocycle. (In de Rham theory, ψ is what Bott and Tu [5] call a ‘global angular form’). One proves that $C\Psi$ is a quasi-isomorphism.

Our map $C\Phi$ can be viewed as a Floer-theoretic analogue of $C\Psi$. Its natural habitat is the theory of ‘Lagrangian matching conditions’ in the sense of [24], alias the ‘quilted’ Floer theory of Wehrheim–Woodward [37]. To prove that Φ is an isomorphism we

show that its low-action part defines a linear isomorphism and then fall back on the algebra of \mathbb{R} -graded complexes. We introduce an algebraic criterion (Lemma 5.4) which is applicable even though there is no monotonicity hypothesis on μ .

To close the discussion, we remark that there are also Gysin-type sequences entirely within Lagrangian Floer homology, and it may well be that a variant of our result can be sharpened to a statement about of Fukaya categories. Assume that we have a rigorous formulation of the monotone Fukaya A_∞ -category \mathcal{F}_N , along the lines of that of the exact Fukaya category in [32]. Fix a monotone Lagrangian $L \in \text{Ob}(\mathcal{F}_N)$. There will then be an A_∞ -module over \mathcal{F}_N [32, Chapter 1], $\text{Cone}(e(V) \smile \cdot)$, sending $L' \in \text{Ob}(\mathcal{F}_N)$ to the chain complex $\text{cone}(e(V) \smile \cdot : \text{CF}^*(L', L) \rightarrow \text{CF}^{*+k+1}(L', L))$; and a second A_∞ -module sending L' to the Floer complex $\text{CF}^*(L' \times \widehat{V}, \widehat{V}^\circ \times L)$ taken in $N \times M_- \times N$, where \widehat{V}° denotes the ‘opposite’ correspondence to \widehat{V} . (For this we need L, L' and \widehat{V}° to be normalised monotone; see Section 3.)

Conjecture 1.16 *Under the hypotheses of Theorem 1, these two A_∞ -modules are quasi-isomorphic.*

Our proof readily adapts to show that $\text{HF}^*(L' \times \widehat{V}, \widehat{V}^\circ \times L) \cong H \text{cone}(e(V) \smile \cdot)$, so the substance of the conjecture lies in the compatibility with the A_∞ -structure.

By the main theorem of [37], $\text{HF}^*(L' \times \widehat{V}, \widehat{V}^\circ \times L) \cong \text{HF}_{M_-}^*(\rho^{-1}L', \rho^{-1}L)$. As the present paper was nearing completion, the author learned that the relation of $\text{HF}^*(\rho^{-1}L, \rho^{-1}L)$ to $\text{HF}^*(L, L)$ has been analysed directly, in the case $k = 1$, by Biran–Khanovsky [4]. They focus on circle-bundles appearing as the boundary of a tubular neighbourhood of symplectic hypersurfaces, and proceed by a direct comparison of the ‘pearl complexes’ computing Floer self-homology. The results appear to be similar to ours where both apply, but the motivations, methods and applications are quite different.

1.5 Navigation

Section 2 reviews the basic Floer-theoretic constructions, as well as some basic facts about symplectic fibrations. Section 3, which is largely to blame for the extravagant length of the paper, explains the framework of ‘Lagrangian matching conditions’. There is more generality here than is really needed, but this framework undoubtedly has further applications. Readers who find such formalism disagreeable are urged to press on to Sections 4 and 5, wherein we prove Theorem 1. In Section 6 we study the ‘borderline’ case, and briefly consider orientations. Section 7 explains the connection

with Seiberg–Witten gauge theory and the monopole Floer homology of connected sums which provided the initial impetus for this paper.

1.6 Acknowledgements

My thanks to Robert Lipshitz, Ciprian Manolescu, Dusa McDuff, Ivan Smith, Katrin Wehrheim and Chris Woodward for stimulating discussions, and to MIT for generous hospitality. I acknowledge financial support during the initial stages of this research from EPSRC research grant EP/C535995/1.

2 Review: fibre bundles, Floer homology

2.1 Locally Hamiltonian fibre bundles

Definition 2.1 A *locally Hamiltonian fibration* (LHF) is a triple (E, p, Ω) consisting of a smooth, proper fibre bundle $p: E \rightarrow B$ equipped with a closed two-form $\Omega \in \Omega^2(E)$ which restricts to each fibre E_b as a non-degenerate (hence symplectic) form.

An LHF carries a canonical connection $A = A(\Omega)$. The horizontal distribution $T^h E$ is given by $T_x^h E = \text{Ann}^\Omega(T_x^\vee E)$: that is, a tangent vector u is horizontal iff $\Omega(u, v) = 0$ for all vertical vectors $v \in T_x^\vee E := \ker Dp$. This connection is symplectic in the sense that $(\mathcal{L}_X \Omega)|_{E_b} = 0$ when X is a horizontal vector field defined in a neighbourhood of E_b . The monodromies of loops based at b are therefore symplectic; so, writing $(M, \omega_M) = (E_b, \Omega|_{E_b})$, and assuming B to be connected, the structure group is reduced to $\text{Aut}(M, \omega_M)$. The monodromies of null-homotopic loops are in fact Hamiltonian automorphisms [18, Chapter 6].

Example 2.2 A fundamental example is the *mapping torus* $(T_\mu, p_\mu, \omega_\mu)$ of an automorphism $\mu \in \text{Aut}(M, \omega)$. Here T_μ is the quotient of $M \times \mathbb{R}$ by the relation $(\mu(x), t) \sim (x, t + 1)$. The projection map $p_\mu: T_\mu \rightarrow S^1$ is an M -bundle. The closed, fibrewise-symplectic two-form ω_μ is the form which pulls back to $M \times \mathbb{R}$ as $\text{pr}_1^* \omega$. Any locally Hamiltonian fibre bundle over S^1 is canonically isomorphic to the mapping torus of its monodromy.

Once we progress to higher-dimensional bases, the connection A is usually not flat. Its curvature F_Ω is a two-form on B with values in the bundle \mathcal{V} of symplectic vector

fields on the fibres, defined via horizontal lifting $X \mapsto X^h$ of vector fields on B in the usual way:

$$(7) \quad F_\Omega \in \Omega^2(B; \mathcal{V}); \quad F_\Omega(X, Y) = [X^h, Y^h] - [X, Y]^h.$$

The function $\kappa_{X,Y} = \Omega(X^h, Y^h)$ on E has the property that

$$(8) \quad d\kappa_{X,Y} = \mathcal{L}_{X^h}\iota(Y^h)\Omega - \mathcal{L}_{Y^h}\iota(X^h)\Omega = \iota([X^h, Y^h])\Omega;$$

hence $\kappa_{X,Y}|_{E_b}$ is the Hamiltonian generator for the vertical vector field $F_\Omega(X, Y)_b$. We can therefore understand F_Ω through the map $(X, Y) \mapsto \kappa_{X,Y}$, which is a two-form on B with values in functions on the fibres. One may in particular speak of *flat* LHF's, meaning those for which $F_\Omega = 0$, or equivalently, those for which the function-valued two-form $(X, Y) \mapsto \kappa_{X,Y}$ on B reduces to an ordinary two-form. The techniques of [18, Chapter 6] readily yield the following.

Lemma 2.3 *A flat LHF (E, p, Ω) over the cylinder $S^1 \times I$, with fibre $(M, \omega) = (E_{(z,0)}, \Omega|_{E_{(z,0)}})$, is equivalent to $(T_\mu \times I, p_\mu \times \text{id}, \text{pr}_1^* \omega_\mu + \beta)$, where μ is the monodromy around $S^1 \times \{0\}$ and β is the pullback of a form on the base.*

If $u: B \rightarrow E$ is a horizontal section, the pullback bundle $u^*T^\vee E \rightarrow B$ inherits a symplectic connection ∇^u , the result of linearizing A_Ω : as noted in [30, Section 2], one has $\nabla_X^u Z = u^*[X^h, \tilde{Z}]$, where X is a vector field on B and \tilde{Z} any section of $T^\vee E$ such that $u^*\tilde{Z} = Z$. The curvature $F(\nabla^u) \in \Omega_B^2(\mathfrak{sp}(u^*T^\vee E))$ is given by

$$(9) \quad F(\nabla^u)(X, Y): Z \mapsto u^*[F_\Omega(X, Y), \tilde{Z}] \in C_B^\infty(u^*T^\vee E).$$

2.2 Floer homology for fixed points and Lagrangian intersections

There are two variants of symplectic Floer homology:

- Floer homology for fixed points of symplectic automorphisms (not necessarily isotopic to the identity); and
- Floer homology for intersections of Lagrangian submanifolds.²

Both are modules over the universal Novikov ring Λ_R of a commutative ring R , as defined at (1.5). Until further notice, we assume that R has characteristic 2.

²The first can be seen as a case of the second, since the Floer homology of an automorphism ϕ is canonically isomorphic to that of the diagonal $\Delta \subset M_- \times M$ with the graph of ϕ , by an isomorphism that essentially goes back to Floer. However, the author does not know a source in the literature that explains this carefully. Note that the graph may not be monotone.

2.2.1 Fixed-point Floer homology

Suppose that (T, p, Ω) is a LHF (2.1) over a closed oriented 1-manifold Z . We suppose the typical fibre N to be compact and either monotone or strongly negative (see (1.8)).

To (T, p, Ω) we attach its Floer homology $\mathrm{HF}(T, \Omega)$. When Z is connected, and $(N, \omega_N) = (E_z, \Omega|_{E_z})$ for a basepoint z , the monodromy along Z is an automorphism $\mu \in \mathrm{Aut}(N, \omega_N)$, and there is a two-form-preserving bundle isomorphism of T with T_μ . We often write $\mathrm{HF}(\mu)$ instead of $\mathrm{HF}(T, \Omega)$.

$\mathrm{HF}(T, \Omega)$ is initially defined when Z is connected and the monodromy μ has non-degenerate fixed points. It is then the homology of a complex $\mathrm{CF}(T, \Omega; J^v)$, where the underlying Λ_R -module $\mathrm{CF}(T, \Omega)$ is the Λ_R -module freely generated by the set $\mathcal{H}(T, \Omega)$ of *horizontal sections* of $T \rightarrow Z$ (a set in bijection with $\mathrm{Fix}(\mu)$). The differential ∂_{J^v} depends on a choice of $J^v \in \mathcal{J}(T^\vee T) = \mathcal{J}(T^\vee T, \Omega)$, the set of Ω -compatible almost complex structure in the vertical tangent bundle $T^\vee T$. The complex also depends on an orientation-preserving diffeomorphism $Z \cong S^1$.

Let $\mathrm{pr}_2: \mathbb{R} \times S^1 \rightarrow S^1$ be the second projection, and form the LHF $(\mathrm{pr}_2^* T, p \circ \mathrm{pr}_2, \mathrm{pr}_2^* \Omega)$ over $\mathbb{R} \times S^1$ (we use the parametrisation $Z \cong S^1$ here). Any $J^v \in \mathcal{J}(T^\vee T)$ extends to a unique almost complex structure \tilde{J}^v in $p^* T$ such that (i) \tilde{J}^v is translation-invariant; (ii) $p \circ \mathrm{pr}_2$ is \tilde{J}^v -holomorphic (where the cylinder is identified with $\mathbb{C}/(z \sim z + i)$ and inherits the latter's complex structure); and (iii) \tilde{J}^v preserves the horizontal distribution. One then considers the moduli space $\mathcal{M}(T; J^v)$ of \tilde{J}^v -holomorphic sections $u: \mathbb{R} \times S^1 \rightarrow \mathrm{pr}_2^* T$ of finite *action*

$$(10) \quad A(u) = \int_{\mathbb{R} \times S^1} u^*(\mathrm{pr}_2^* \Omega).$$

We give it the $L^p_{1,loc}$ topology, for some $p > 2$. When T is non-degenerate, any $u \in \mathcal{M}(T; J^v)$, thought of as a path of sections $u(s, \cdot)$ of T , converges uniformly as $s \rightarrow \pm\infty$ to horizontal sections $\nu_\pm \in \mathcal{H}(T, \Omega)$ [26]. Hence $\mathcal{M}(T; J^v)$ is partitioned into components $\mathcal{M}(T; \nu_-, \nu_+; J^v)$. The *Conley–Zehnder index* $\mathrm{ind}: \mathcal{M}(T; \nu_-, \nu_+; J^v) \rightarrow \mathbb{Z}$, which coincides with the Fredholm index of the linearisation of the Cauchy–Riemann equation $\tilde{J}^v \circ Du = Du \circ i$, further partitions $\mathcal{M}(T; \nu_-, \nu_+; J^v)$ into subsets $\mathcal{M}(T; \nu_-, \nu_+; J^v)^q := \mathrm{ind}^{-1}(q)$.

Definition 2.4 We define a Baire subset $\mathcal{J}^{\mathrm{reg}}(T^\vee T) \subset \mathcal{J}(T^\vee T)$ of *regular* complex structures as follows. When N is monotone, $J^v \in \mathcal{J}^{\mathrm{reg}}(T^\vee T)$ if the moduli spaces $\mathcal{M}(T; \nu_-, \nu_+; J^v)^q$ are all transversely cut out (in particular, empty when $q \leq 0$). In the strongly negative case, we require additionally that there should be no non-constant J^v -holomorphic spheres in the fibres of T .

Remark Baire subsets of Fréchet spaces need not be dense. To get around this, we can make an initial choice of $J \in \mathcal{J}(T^\vee T)$, and then define regularity among those $J' \in \mathcal{J}(T^\vee T)$ which are close to J in Floer's C_ϵ^∞ norm [15]; the norm depends on a choice of rapidly decreasing sequence $\epsilon = (\epsilon_m)_{m=1,2,\dots}$ and on a metric on $T^\vee T$. This detail will be glossed over from now on.

For a proof that transversality of the moduli spaces is generic see, e.g., [15]. To understand why non-existence of J^\vee -spheres is generic in the strongly negative setting, look at the moduli space $\mathcal{N}^s(T)$ of non-constant, simple, parametrised J^\vee -spheres $\beta: S^2 \rightarrow T$, mapping to some (variable) fibre of p . The real, fibrewise index $l = 2 \int_{S^2} \beta^* c_1(TN) + \dim(N)$ breaks this moduli space into components $\mathcal{N}^s(T)^l$. Impose the (generic) condition on J^\vee that each $\mathcal{N}^s(T)^l$ is transversely cut out. The quotient of $\mathcal{N}^s(T)^l$ by $\text{Aut}(S^2) = \text{PSL}(2, \mathbb{C})$ is then a smooth manifold of dimension $l + 1 - \dim \text{Aut}(S^2) = l - 5$ [19]. But when $c_1^{\min}(N) \geq n - 2$, this dimension is ≤ -1 for all β with $\langle \beta, [\omega] \rangle \geq 0$, and so $\mathcal{N}^s(T)^l$ is empty for all l . Since every sphere factors through a simple one, there are then no holomorphic spheres in the fibres of T .

We shall need to consider the subset $\mathcal{M}(T; J^\vee)_q^r \subset \mathcal{M}(T; J^\vee)^q$ of sections u whose action $A(u)$ equals r . We shall also need to divide out by the \mathbb{R} -action, by translation, on the moduli spaces. For any $J^\vee \in \mathcal{J}^{\text{reg}}(T^\vee T)$, any $\nu_-, \nu_+ \in \mathcal{H}(T, \Omega)$, and any $r \geq 0$, the space $\mathcal{M}(T; \nu_-, \nu_+; J^\vee)_1^r / \mathbb{R}$ is compact and 0-dimensional. Say it consists of $n(\nu_-, \nu_+)_r$ points, modulo 2. Compactness of $\mathcal{M}(T; \nu_-, \nu_+; J^\vee)_1^r / \mathbb{R}$ follows from Gromov–Floer compactness and the regularity assumptions, which forbid bubbling in the index 1 case. The differential ∂_{J^\vee} is defined by

$$(11) \quad \partial_{J^\vee} \langle \nu_- \rangle = \sum_{\nu_+} \left(\sum_{r \in \mathbb{R}} n(\nu_-, \nu_+)_r t^r \right) \langle \nu_+ \rangle.$$

One has $\partial_{J^\vee}^2 = 0$, so that this is a well-formed complex $\text{CF}(T, \Omega; J^\vee)$. One then uses ‘continuation maps’ to prove the invariance of $\text{HF}(T, \Omega)$ under change of almost complex structure, of parametrisation of Z , and under changes $\Omega \rightsquigarrow \Omega + d\alpha$ where $d\alpha$ vanishes on the fibres. Since the monodromy can always be made non-degenerate by such perturbations, the definition can be relaxed to allow degenerate (T, Ω) . See [15] for proofs (trivial changes must be made to generalise from the case of Hamiltonian symplectomorphisms).

Finally, when Z is disconnected, one takes the tensor product of the chain complexes attached to the components of Z .

2.2.2 Lagrangian intersection Floer homology

Suppose that L_0 and L_1 are monotone Lagrangian submanifolds in an (M, ω_M) which is closed or has convex contact-type boundary. Assume that $\min(m_{L_0}^{\min}, m_{L_1}^{\min}) \geq 3$. One then has a Lagrangian Floer homology module $\mathrm{HF}_*(L_0, L_1) = H(\mathrm{CF}_*(L_0, L_1; J))$ over Λ_R [21].

The chain complex $\mathrm{CF}(L_0, L_1; J)$ is defined on the free Λ_R -module $\mathrm{CF}(L_0, L_1)$ generated by $L_0 \cap L_1$ (here one first moves L_1 by a Hamiltonian diffeomorphism to ensure the intersection is transverse). The definition of the differential δ_J invokes an element $J = \{J_t\}_{t \in [0,1]}$ of the space $\mathcal{PJ}(M)$ of smooth paths $[0, 1] \rightarrow \mathcal{J}(M)$ in the space of ω_M -compatible almost complex structures. One forms a moduli space $\mathcal{M}(L_0, L_1; J)$ of $L_{1,loc}^p$ maps $u: \mathbb{R} \times [0, 1] \rightarrow M$ with $u(\mathbb{R} \times \{i\}) \subset L_i$ ($i = 1, 2$), of finite action $A(u) = \int_{[0,1] \times \mathbb{R}} u^* \omega_M$, satisfying the Cauchy–Riemann equation $J_t \circ D_{(s,t)} u = D_{(s,t)} u \circ i$. The subset $\mathcal{PJ}^{\mathrm{reg}}(M) \subset \mathcal{PJ}(M)$ is defined by the requirement that the components of the quotient space $\mathcal{M}(L_0, L_1; J)/\mathbb{R}$ of formal dimension ≤ 1 are cut out transversely.

For $J \in \mathcal{PJ}^{\mathrm{reg}}(M)$, monotonicity and the Maslov index assumption allow one to rule out the bubbles and boundary bubbles which potentially lead to the failure of compactness for moduli spaces of broken trajectories. This is done by an index argument (applied twice: once in proving that $\delta_J \circ \delta_J = 0$, and once in defining the continuation maps which prove invariance). Thus $\mathrm{HF}(L_0, L_1)$ is well-defined and independent of J and of Hamiltonian perturbations of L_0 and L_1 . This is also true when $m_{L_0}^{\min} = 2$ and $L_1 = \phi(L_0)$ for some $\phi \in \mathrm{Aut}(M, \omega)$ [21, Addendum] (discs bubble off in cancelling pairs).

Floer homology is well-defined and invariant also when (M, ω) and the two Lagrangians L_0 and L_1 are all strongly negative (compare [17]). The definition of $\mathcal{PJ}^{\mathrm{reg}}(M)$ is slightly more elaborate than in the monotone case: the additional requirements are that the moduli spaces of simple J_t -spheres, somewhere-injective J_0 -discs attached to L_0 , and somewhere-injective J_1 -discs attached to L_1 are all cut out transversely. This ensures that these moduli spaces are negative-dimensional manifolds, hence empty. The possibility of *nowhere*-injective discs is ruled out by a lemma of Kwon–Oh and Lazzarini [16, 17] to the effect that the image of any pseudo-holomorphic disc attached to a Lagrangian contains that of a somewhere-injective disc.

2.3 Open manifolds

In Theorem 1 we allow (M, ω) to have a convex contact-type boundary: thus there is a non-vanishing, outward-pointing vector field X defined near ∂M , satisfying the

Liouville condition $L_X\omega = \omega$. For the purposes of pseudo-holomorphic theory on such manifolds, one imposes a convexity condition on almost complex structures (see [30, Section 2.1]) which is understood as part of the notion of a ‘compatible’ almost complex structure. The adaptations described in [30] do not interact at all with the proof of the theorem in Sections 4 and 5; therefore, to avoid irrelevant complication, in our exposition we shall only consider closed manifolds.

3 Floer homology and Lagrangian matching conditions

3.1 A cobordism category for Lagrangian correspondences

Floer homology groups for Lagrangian correspondences fit into a sort of TQFT, or sigma-model, of a non-standard kind, which we now outline. It could be more accurately described as a $(1 + 1)$ -dimensional topological field theory coupled to locally Hamiltonian fibrations and subject to ‘Lagrangian matching conditions’. It is under systematic development by Wehrheim–Woodward [37], though its existence was observed independently by the author. I have decided, after some hesitation, to outline the general theory, since it is an appropriate setting for the symplectic Gysin sequence. The theory as we shall describe it differs from that of [37] only in some rather minor aspects.

We first describe the objects in a cobordism category, \mathcal{C} (the morphisms will be described later).

Definition 3.1 A *cycle of Lagrangians*, of length $l \geq 1$, is a sequence of symplectic manifolds $M = (M_1, \dots, M_l)$ together with a cycle

$$L = (L_{1,2}, L_{2,3}, \dots, L_{l-1,l}, L_{l,1})$$

of Lagrangian correspondences³

$$L_{i,i+1} \subset (M_i)_- \times M_{i+1}, \quad i \in \mathbb{Z}/l.$$

Definition 3.2 An object O of \mathcal{C} comprises a collection $\{(M_\alpha, L_\alpha)\}_{\alpha \in A}$ of cycles of Lagrangians (M_α, L_α) , indexed by a finite set A , and a collection $\{(M_\beta, \omega_\beta, \mu_\beta)\}_{\beta \in B}$ of compact symplectic manifolds (M_β, ω_β) with symplectic automorphisms $\mu_\beta \in \text{Aut}(M_\beta, \omega_\beta)$, indexed by another finite set B .

³A Lagrangian correspondence between M_1 and M_2 is, by definition, a Lagrangian submanifold of $(M_1)_- \times M_2$.




Figure 1: Schematic of a mapping torus (left) and of a cycle of Lagrangians of length 4 (right).

To give a symplectic automorphism $\mu \in \text{Aut}(M, \omega)$ is (essentially) equivalent to giving an oriented 1-manifold $Z \cong S^1$ with a marked point z , and an LHF (E, π, Ω) over Z : given μ one forms its mapping torus; given (E, π, Ω) and $z \in Z$ one takes the monodromy on E_z .

Cycles of Lagrangians can also be described in fibre-bundle language. Define a *cycle of LHF's* in terms of the following data: an oriented 1-manifold $Z \cong S^1$ with a finite set of marked points enumerated in cyclic order as (z_1, \dots, z_l) dividing Z into intervals $[z_i, z_{i+1}]$; and, for each $i \in \mathbb{Z}/l$, an LHF (E_i, π_i, Ω_i) over $[z_{i-1}, z_i]$, together with a Lagrangian

$$\Lambda_{i,i+1} \subset \pi_i^{-1}(z_i)_- \times \pi_{i+1}^{-1}(z_i).$$

Write M_i for $\pi_i^{-1}(z_{i-1})$. Parallel transport $m_i: M_i \rightarrow \pi_i^{-1}(z_i)$ trivialises (E_i, π_i, Ω_i) . One has Lagrangian correspondences $L_{i,i+1} = (\text{id} \times m_i^{-1})\Lambda_{i,i+1} \subset (M_i)_- \times \pi_i^{-1}(z_i)$. Thus one obtains a cycle of Lagrangians $(M_i; L_{i,i+1})_{i \in \mathbb{Z}/l}$. Conversely, given a cycle of Lagrangians, one can build a cycle of (trivial) LHF's. These constructions are mutually inverse (though only up to isomorphism when one starts with a cycle of LHF's).

For schematic pictures of a mapping torus and a cycle of Lagrangians, see Figure 1.

3.1.1 Non-degeneracy

A cycle of Lagrangians $L_{i,i+1} \subset (M_i)_- \times M_{i+1}$ ($i \in \mathbb{Z}/l$) is called *non-degenerate* when the product $\mathbb{L} := \prod_{i \in \mathbb{Z}/l} L_{i,i+1}$ (a Lagrangian in $\prod_{i \in \mathbb{Z}/l} (M_i)_- \times M_{i+1}$) meets the multi-diagonal $\Delta_l = \{(x_1, x_2, x_2, \dots, x_l, x_l, x_1)\}$ (another Lagrangian) transversely.

Lemma 3.3 $\mathbb{L} = \prod_{i \in \mathbb{Z}/l} L_{i,i+1}$ may be moved by a Hamiltonian isotopy to a new Lagrangian, close to the old one in Floer's C_ϵ^∞ -norm, but meeting Δ_l transversely.

Moreover, the Hamiltonian can be taken to be a product of Hamiltonian isotopies $\text{pr}_2^*(H_{i,t})$ in the factors $(M_i)_- \times M_{i+1}$.

Proof Consider the map

$$q: \prod_i C^\infty([0, 1] \times M_{i+1}) \times \mathbb{L} \rightarrow \mathbb{M}$$

given by

$$(H_1, \dots, H_l; (x_1, y_1), \dots, (x_l, y_l)) \mapsto ((x_1, \phi_{H_1}(y_1)), \dots, (x_l, \phi_{H_l}(y_l))).$$

Observe that $\text{im}(Dq)$ is the subspace $\bigoplus_i T_{\phi_{H_i}(y_i)} M_{i+1} \subset \bigoplus_i T_{x_i} M_i \oplus T_{\phi_{H_i}(y_i)} M_{i+1}$. Thus q is transverse to Δ_l . Form the fibre product

$$Z = \left(\prod_i C^\infty([0, 1] \times M_{i+1}) \times \mathbb{L} \right) \times_q \Delta_l,$$

and note that transversality of $(\prod \text{id}_{M_i} \times \phi_{H_i})(\mathbb{L})$ with Δ_l is equivalent to (H_1, \dots, H_l) being a regular value of the projection $Z \rightarrow \prod_i C^\infty([0, 1] \times M_{i+1})$. The projection is Fredholm, so the Sard–Smale theorem gives the density of the regular values. \square

There is a corresponding notion of non-degeneracy for cycles of LHF's. A *cycle of Hamiltonian chords* $(\sigma_1, \dots, \sigma_l)$ is a sequence of Ω_i -horizontal sections $\sigma_i: [z_{i-1}, z_i] \rightarrow E_i$ where $(\sigma_{i-1}(z_i), \sigma_i(z_i)) \in L_{i-1,i}$ for each $i \in \mathbb{Z}/l$. Cycles of Hamiltonian chords are in natural bijection with intersections $\mathbb{L} \cap \Delta_l$ for the corresponding cycle of Lagrangians, and one says that the cycle of LHF's is non-degenerate when the intersection is transverse. It follows from the genericity property just discussed that non-degeneracy can be achieved by making a small adjustment to the forms Ω_i , replacing each by $\Omega_i + d(K_i dt)$ where t is the circle coordinate.

3.2 Floer homology from cycles of correspondences

Before turning to the morphisms in the category \mathcal{C} , we explain how to attach homology groups to objects.

Definition 3.4 A compact symplectic manifold (M, ω) is *normalised monotone* if $\int_{S^2} f^* \omega = 2 \langle f^* c_1(TM, \omega), [S^2] \rangle$ for all smooth maps $f: S^2 \rightarrow M$. A compact Lagrangian $L \subset M$ is *normalised monotone* if $\int_D f^* \omega = m_L(f)$ for all $f: (D^2, \partial D^2) \rightarrow (M, L)$.

Notice that the product of normalised monotone symplectic manifolds is normalised monotone. From a cycle of normalised monotone Lagrangian correspondences $L_{i,i+1} \subset (M_i)_- \times M_{i+1}$ between normalised monotone manifolds $\{M_i\}_{i \in \mathbb{Z}/l}$, we form the product

$$(12) \quad \mathbb{M} = \prod_{i \in \mathbb{Z}/l} (M_i)_- \times M_{i+1} = \prod_{i \in \mathbb{Z}/l} M_i \times (M_i)_-$$

(also normalised monotone) and two normalised monotone Lagrangians inside it: the product $\mathbb{L} = \prod_{i \in \mathbb{Z}/l} L_{i,i+1}$, and the multi-diagonal Δ_l , consisting of cycles whose components in M_i and in $(M_i)_-$ coincide for each $i \in \mathbb{Z}/l$.

For each $\mathbb{J} \in \mathcal{J}^{\text{reg}}(\mathbb{M}; \mathbb{L}, \Delta_l)$ one has a Floer complex

$$(13) \quad \text{CF}(\mathbb{L}, \Delta_l; \mathbb{J}),$$

whose chain-homotopy type is independent of \mathbb{J} . We wish to show that regularity can be attained when \mathbb{J} is drawn from a smaller pool of almost complex structures.

Recall that $\mathcal{J}(M_i)$ is the space of compatible almost complex structures on M_i . Put $\mathcal{J}(M_1, \dots, M_l) = \prod_{i \in \mathbb{Z}/l} \mathcal{J}(M_i)$. The *pleating map* pl ,

$$(14) \quad \text{pl}: \mathcal{P}\mathcal{J}(M_1, \dots, M_l) \rightarrow \mathcal{P}\mathcal{J}(\mathbb{M})$$

sends the path $t \mapsto (J_1(t), \dots, J_l(t))$ to the path

$$t \mapsto (J^1(t), -J^1(1-t); \dots; J^l(t), -J^l(1-t)).$$

(Here we think of \mathbb{M} as $\prod_{i \in \mathbb{Z}/l} M_i \times (M_i)_-$). Let $\mathcal{P}\mathcal{J}^{\text{reg}}$ denote those products of t -dependent almost complex structures which induce regular almost complex structures for the pair (\mathbb{L}, Δ_l) :

$$(15) \quad \mathcal{P}\mathcal{J}^{\text{reg}}(M_1, \dots, M_l; \mathbb{L}, \Delta_l) = \text{pl}^{-1}(\mathcal{P}\mathcal{J}^{\text{reg}}(\mathbb{M}; \mathbb{L}, \Delta_l)).$$

Lemma 3.5 $\mathcal{P}\mathcal{J}^{\text{reg}}(M_1, \dots, M_l; \mathbb{L}, \Delta_l)$ is a Baire subset of $\mathcal{P}\mathcal{J}(M_1, \dots, M_l)$.

Proof This is a consequence of the Sard–Smale theorem, and only the last stage of the proof diverges from the standard arguments found in, e.g., [30, Lemma 2.4]. The last stage is the following: let $u: [0, 1] \times \mathbb{R} \rightarrow \mathbb{M}$ be a non-constant, finite-action (j, \mathbb{J}) -holomorphic map satisfying the boundary conditions $u(\{0\} \times \mathbb{R}) \subset \mathbb{L}$ and $u(\{1\} \times \mathbb{R}) \subset \Delta_l$. Let $F \rightarrow [0, 1] \times \mathbb{R}$ be the bundle dual to $\Lambda^{0,1}(u^*T\mathbb{M})$, and suppose that η is a section of F , of class $L^{p/(p-1)}$, such that $D_{u, \mathbb{J}}^* \eta = 0$ on $(0, 1) \times \mathbb{R}$ and such that

$$\int_{[0,1] \times \mathbb{R}} \langle \eta, \mathbb{Y} \circ Du \circ j \rangle = 0$$

for all $\mathbb{Y} \in T_J \mathcal{J}(M_1, \dots, M_l)$. Concretely,

$$\mathbb{Y}(t) = (Y_1(t), -Y_1(1-t); \dots; Y_l(t), -Y_l(1-t)),$$

where $Y_i(t) \in \text{End}(TM_i)$ satisfies $J_i(t) \circ Y_i(t) + Y_i(t) \circ j = 0$. We must prove that $\eta = 0$.

By elliptic bootstrapping, η is actually continuous. The set of points at which $Du = 0$ is discrete [11], so if η is not everywhere zero, there is a point $(t, s) \in [0, 1] \times \mathbb{R}$ where $\eta(t, s) \neq 0$ and $Du(x, y) \neq 0$. We may assume $t \neq 1-t$. Let $\alpha = (Du \circ j)(t, s) \in \text{Hom}(\mathbb{R}^2, T_{u(t,s)}\mathbb{M})$; it has components $(\beta_1, \gamma_1; \dots; \beta_l, \gamma_l)$. Say $(\beta_i, \gamma_i) \neq 0$; then there exist $Y_i(t) \in \text{End}(TM_i)$ satisfying $J_i(t) \circ Y_i(t) + Y_i(t) \circ j$ and $Y_i(1-t) \in \text{End}(TM_i)$ satisfying $J_i(1-t) \circ Y_i(1-t) + Y_i(1-t) \circ j$, such that also

$$Y_i(t)(\beta_i) \neq Y_i(1-t)(\gamma_i).$$

Define $\mathbb{Y}(t)$ by putting $(Y_i(t), -Y_i(1-t))$ in the i th slot, and $(Y_j(t), Y_j(1-t)) = (0, 0)$ for $j \neq i$. Then $\langle \eta(s, t), \mathbb{Y}(t) \circ \alpha \rangle \neq 0$. Extend $\mathbb{Y}(t)$ arbitrarily to a $\tilde{\mathbb{Y}} \in T_J \mathcal{J}(M_1, \dots, M_l)$, then replace $\tilde{\mathbb{Y}}$ by $\mathbb{Y} = \chi \tilde{\mathbb{Y}}$ where χ is a bump function supported near (s, t) . When the support of χ is small enough, one will then have $\int_{[0,1] \times \mathbb{R}} \langle \eta, \mathbb{Y} \circ Du \circ j \rangle \neq 0$, contrary to assumption. Hence $\eta \equiv 0$.

Strictly, what this argument proves is that (for a Baire subset of \mathbb{J} 's), all non-constant Floer-theoretic trajectories are cut out transversely. In the strongly negative case, there are additional transversality conditions on bubbles. However, these can be handled by almost identical arguments to those for trajectories. \square

Remark In forming the complex (13), we shall generally choose \mathbb{J} from the space $\mathcal{P}\mathcal{J}^{\text{reg}}(M_1, \dots, M_l; \mathbb{L}, \Delta_l)$. Thus we assemble \mathbb{J} from t -dependent almost complex structures on l (rather than $2l$) symplectic manifolds.

Proposition 3.6 *When l is even, and $\mathbb{J} = \text{pl}(J_1, \dots, J_l) \in \mathcal{J}^{\text{reg}}(M_1, \dots, M_l; \mathbb{L}, \Delta_l)$, there is a canonical isomorphism of chain complexes*

$$\text{CF}(\mathbb{L}, \Delta_l; \mathbb{J}) \cong \text{CF}(L_{1,2} \times L_{3,4} \times \dots \times L_{l-1,l}, L_{2,3} \times L_{4,5} \times \dots, L_{l,1}; J_1, \dots, J_l),$$

where the latter complex is taken in $M_1 \times (M_2)_- \times \dots \times M_{l-1} \times (M_l)_-$. In particular, when $l = 2$, we have

$$\text{CF}(L_{1,2} \times L_{2,1}, \Delta_2) \cong \text{CF}(L_{1,2}, L_{2,1}).$$

Proof A generator for $\text{CF}(\mathbb{L}, \Delta_l; \mathbb{J})$ is an intersection point $((x_1, x_2), (x_2, x_3), \dots, (x_l, x_1))$ with $(x_i, x_{i+1}) \in L_{i,i+1}$; we map this to the generator (x_1, x_2, \dots, x_l) for the complex on the right and so obtain a linear isomorphism between the two complexes. Checking that this is a chain map is left to the reader. \square

An object O of \mathcal{C} is, according to Definition 3.2, a finite collection (indexed by some set A) each of whose elements is a sequence of symplectic manifolds $(M_1^\alpha, \dots, M_{l_\alpha}^\alpha)$ and a cycle of Lagrangian correspondences $L_{i,i+1}^\alpha$, together with a finite collection (indexed by B) of symplectic manifolds N_β equipped with automorphisms μ_β .

Given a monotone (or strongly negative) N and a non-degenerate automorphism $\mu \in \text{Aut}(N)$, we have a family of chain complexes $\text{CF}(\mu; J^\nu)$ indexed by $J^\nu \in \mathcal{J}^{\text{reg}}(T^\vee T(\mu))$.

We associate with O the homology

$$(16) \quad \text{HF}(O) = H \left(\bigotimes_{\alpha \in A} \text{CF}_*(\mathbb{L}_\alpha, \Delta_{l_\alpha}) \otimes \bigotimes_{\beta \in B} \text{CF}(\mu_\beta) \right),$$

which is independent, up to canonical and coherent isomorphisms, of the additional choices of $\mathbb{J}_\alpha \in \mathcal{J}^{\text{reg}}(\mathbb{L}_\alpha, \Delta_{l_\alpha})$ and $J_\beta^\nu \in \mathcal{J}(T^\vee T(\mu_\beta))$.

3.2.1 Morphisms

A morphism in \mathcal{C} is an isomorphism class of ‘cobordisms’, where a cobordism is a *matched collection of fibrations*. The latter is defined to be a formidable assembly of data

$$(17) \quad \mathcal{F} = (S, \Gamma; E, \pi, \Omega; \mathcal{Q}; \{\zeta_\alpha\}_{\alpha \in \pi_0(\partial S)}),$$

consisting of the following components:

- (i) A quilted surface (S, Γ) ;
- (ii) an LHF (E, π, Ω) over the surface S_Γ obtained by cutting along Γ ;
- (iii) a Lagrangian matching condition $\mathcal{Q} \rightarrow \Gamma$ for E ;
- (iv) distinguished points ζ^α in the boundary.

Here are the definitions.

Component (i):

Definition 3.7 A *quilted surface* (Wehrheim–Woodward’s term) (S, Γ) consists of a compact oriented surface S , possibly with boundary, and a smoothly embedded 1-manifold with boundary Γ , intersecting ∂S transversely, with $\partial\Gamma \subset \partial S$.




Figure 2: Above: a quilted cylinder (S, Γ) . Below: the corresponding surface S_Γ with the three parts of its boundary labelled.

We denote by S_Γ the compact surface with boundary and corners obtained by cutting S open along Γ , i.e. by adjoining boundary arcs and circles to compactify $S \setminus \Gamma$. The boundary of S_Γ consists of two copies of Γ which we denote Γ^+ and Γ^- (they may be taken in either order), and a union of closed arcs and circles T coming from ∂S and meeting $\Gamma^+ \cup \Gamma^-$ at the corners. Thus

$$\partial S_\Gamma = \Gamma^+ \cup \Gamma^- \cup T.$$

Component (ii): this is an LHF $(\pi: E \rightarrow S_\Gamma, \Omega)$. It is allowed to have differing dimensions over different components of S_Γ . However, we insist that the fibration is flat over a neighbourhood of T . We shall be particularly interested in its restrictions $E|_{\Gamma^\pm} \rightarrow \Gamma^\pm$, and in the closed two-forms Ω^\pm on $E|_{\Gamma^\pm}$ obtained by restricting Ω .

It is important to note that Γ^+ is naturally identified with Γ (and that the same goes for Γ^-). Thus we may regard $E|_{\Gamma^\pm}$ as LHF's over Γ . Form the fibre product

$$(18) \quad E_\Gamma := (E|_{\Gamma^+}) \times_\Gamma (E|_{\Gamma^-}).$$

It becomes an LHF over Γ when we endow it with the two-form Ω_Γ obtained as the restriction of the form $(-\Omega^+) \oplus \Omega^-$ on $(E|_{\Gamma^+}) \times (E|_{\Gamma^-})$ to its subspace E_Γ . The minus sign is crucial (though to which factor it is attached is not).

Component (iii):

Definition 3.8 A Lagrangian matching condition $Q \rightarrow \Gamma$ for E is a sub-fibre bundle $Q \subset E_\Gamma$, proper over Γ , such that (i) $\Omega_\Gamma|_Q = 0$, and (ii) the fibres of Q have half the dimension of those of E_Γ .

We can equally regard Q as a sub-fibre-bundle of $(E|_{\Gamma^-}) \times_{\Gamma} (E|_{\Gamma^+})$; the ordering of Γ^+ and Γ^- is immaterial.

Component (iv): a collection of reference points $\{\zeta^\alpha\}_{\alpha \in \pi_0(\partial S)}$, one in each component $(\partial S)^\alpha$ of ∂S . We require that $\zeta^\alpha \in \partial\Gamma$ when $\partial\Gamma \cap (\partial S)^\alpha \neq \emptyset$.

This concludes the definition of matched collections of fibrations.

Now consider what all this gives us over the original boundary ∂S . For any component $Z = (\partial S)^\alpha$ of ∂S , the matched collection of fibrations gives by restriction a cycle of LHF's. The marked points are those of the finite set $\partial\Gamma \cap Z$, enumerated in cyclic order, consistently with the boundary orientation, as $(z_1^\alpha, \dots, z_{l_\alpha}^\alpha)$, starting with $z_1^\alpha = \zeta^\alpha$. When $\partial\Gamma \cap Z = \emptyset$, we just take the one-element sequence given by z^α . As described above, a cycle of LHF's can be trivialised so as to give a cycle of Lagrangians.

Definition 3.9 A *cobordism* from (M_-, L_-) to (M_+, L_+) is defined to be a matched collection of fibrations, together with a symplectic identification of the cycles of Lagrangians obtained from its boundary with $\overline{M}_- \amalg M_+$. Here \overline{M}_- is the object obtained from M_- by reversing the direction of each of its constituent sequences; thus (M_1, \dots, M_l) becomes (M_l, \dots, M_1) , and the Lagrangian correspondences are obtained in the obvious way (when $l = 0$, μ changes to μ^{-1}). A morphism in the category \mathcal{C} is defined as an isomorphism class of cobordisms.

The definition of ‘isomorphism’ here is the obvious one, involving diffeomorphisms of the base surfaces and two-form-preserving bundle-maps; it is left to the reader to elaborate.

Morphisms are composed by concatenation of cobordisms. The picture to have in mind is that one can *cut* the quilted surface (S, Γ) (and the fibrations lying above it) along a collection of circles transverse to Γ ; sewing is just the inverse operation. For the sewing of fibrations, it is convenient to be able to assume that the LHF's are flat, hence (by Lemma 2.3) locally trivial near the boundaries being glued (compare [30, Section 1]). The formalities are again left to the reader.

3.3 Cobordism-maps

We now turn to the construction of maps on Floer homology from matched collections of fibrations.

3.3.1 Cylindrical ends; almost complex structures

Let \mathcal{F} be a matched collection of fibrations. We elongate the base surface S to a surface \widehat{S} with cylindrical ends. To do this, choose for each boundary component $(\partial S)_\alpha$ a boundary collar $e_\alpha: (-1, 0] \times S^1 \hookrightarrow S$ such that $e_\alpha^{-1}(\Gamma)$ is a union of line segments $(-1, 0] \times \{t\}$. Then define

$$(19) \quad \widehat{S} = S \cup \bigcup_{\alpha} (-1, \infty) \times S^1$$

where the α th copy of $(-1, \infty) \times S^1$ is glued on along the region $(-1, 0] \times S^1$ via the collar e_α . Extend Γ linearly to a 1-manifold $\widehat{\Gamma} \subset \widehat{S}$: that is, if $e_\alpha^* \Gamma$ contains a segment $(-1, 0] \times \{t\}$, one extends it to $(-1, \infty) \times \{t\}$. Let \widehat{S}_Γ be the surface with boundary obtained by cutting \widehat{S} along $\widehat{\Gamma}$.

We may choose e so that the LHF $(E \rightarrow S_\Gamma, \Omega)$ is flat near T (recall the flatness assumption in the definition). It then extends uniquely to an LHF $(\widehat{E} \rightarrow \widehat{S}_\Gamma, \widehat{\Omega})$ which is flat over the ends. The Lagrangian matching condition Q extends to a Lagrangian matching condition \widehat{Q} over $\widehat{\Gamma}$ in a unique way; indeed, the requirement that it be isotropic implies that the extension must be obtained by symplectic parallel transport along the ends of $\widehat{\Gamma}$.

Definition 3.10 (a) Given an LHF (E, π, Ω) over a Riemann surface (S, j) , an almost complex structure J on E is *partly compatible* if $J \circ D\pi = D\pi \circ j$, and the restriction of J to each fibre E_x is compatible with the symplectic form $\Omega|_{E_x}$.

(b) A partly compatible J is *fully compatible* if it preserves the horizontal distribution.

(c) Suppose S has a cylindrical end, holomorphically modelled on $(-\infty, -T] \times S^1$ or $[T, \infty) \times S^1$. We say J is *adapted* if it is partly compatible everywhere, fully compatible over the end, and translation-invariant (after shortening the end by increasing T).

(d) When the cylindrical ends of \widehat{S} have seams $\widehat{\Gamma} = (-\infty, -T) \times \{z_1, \dots, z_l\} \subset (-\infty, -T) \times S^1$, we choose j to be *standard* so that the end is modelled holomorphically on $(-\infty, -T) \times S^1$ with seams $(-\infty, -T) \times \{e^{2\pi i/l}, e^{4\pi i/l}, \dots, e^{2\pi i(l-1)/l}\}$, i.e., the z_k are equally spaced. Given an LHF over the resulting Riemann surface \widehat{S}_Γ , we consider almost complex structures J which are adapted in the same sense as (c).

3.3.2 Holomorphic sections

The ends of the surface \widehat{S} are cylindrical, with seams $\widehat{\Gamma}$. Fix a complex structure j on \widehat{S} (inducing \hat{j} on \widehat{S}_Γ) which is standard over the ends, as in Definition 3.10, and consider

the space $\mathcal{J}(\widehat{E}_\Gamma) = \mathcal{J}(\widehat{E}_\Gamma, j)$ of j -adapted almost complex structures in \widehat{E}_Γ . Over each end e , J is eventually translation invariant and preserves the horizontal distribution, and is there determined by its restriction $J^v(e)$ to the vertical tangent bundle. We write $\mathcal{J}(\widehat{E}_\Gamma; \{J_e^v\})$ for the subspace of $\mathcal{J}(\widehat{E}_\Gamma)$ of almost complex structures limiting to the vertical complex structures J_e^v over the ends. The J_e^v are the complex structures used in the definition of the Floer complexes, and so must be chosen *regular* (cf. Definition 2.4).

Given a section u of $\widehat{E} \rightarrow \widehat{S}_\Gamma$ one obtains, by restriction to the boundary, a section u_Γ of \widehat{E}_Γ (the completed version of E_Γ). At a point $x \in \Gamma$, the value of $u_\Gamma(x)$ is the pair $(u(x_+), u(x_-))$, where x_+ and x_- are the points of Γ^+ and Γ^- that correspond to x under the identifications $\Gamma^+ \cong \Gamma \cong \Gamma^-$.

For a fixed $p > 2$, the configuration space \mathcal{B}^p associated with our non-degenerate matched collection of fibrations is the disjoint union of subspaces $\mathcal{B}^p(\nu)$ indexed by cycles of Hamiltonian chords ν . The space $\mathcal{B}^p(\nu)$ consists of the sections u of $\widehat{E} \rightarrow \widehat{S}_\Gamma$ of class $L_{1,\text{loc}}^p$ with the properties (i) that u_Γ is a section of $Q \subset \widehat{E}_\Gamma$, and (ii) the function $e_\alpha^* u(s, t) - \nu_\alpha(t)$ is an L_1^p -function on the cylinder for each cylindrical end e_α . The L_1^p topology makes this space into a Banach manifold. The *action* of $u \in \mathcal{B}^p(\nu)$ is

$$(20) \quad A(u) = \int_{\widehat{S}_\Gamma} u^* \Omega,$$

when the integral exists. Given an adapted almost complex structure J , the moduli space of J -holomorphic sections is the subspace

$$(21) \quad \mathcal{Z}(\nu; J) := \{u \in \mathcal{B}(\nu) : J \circ Du - Du \circ j = 0, A(u) < \infty\} \subset \mathcal{B}^p(\nu).$$

Thus $\mathcal{Z}(\nu; J) = \bar{\partial}_J^{-1}(0)$ where $\bar{\partial}_J$ is the section $\bar{\partial}_J = J \circ D - D \circ j$ of the natural Banach vector bundle $\mathcal{V}^p \rightarrow \mathcal{B}^p(\nu)$ whose fibre at u is the subspace

$$(22) \quad (\mathcal{V}^p)_u \subset L^p(\Lambda_{\widehat{S}_\Gamma}^{0,1}(u^* T\widehat{E}))$$

consisting of those \mathbb{C} -antilinear homomorphisms $\alpha : T\widehat{S}_\Gamma \rightarrow u^* T\widehat{E}$ whose restrictions to Γ satisfy the linearised matching condition

$$(23) \quad (\alpha_{x_+}(\xi), \alpha_{x_-}(\xi)) \in T_{((u(x_+), u(x_-)))} Q \text{ for all } \xi \in T_x \widehat{\Gamma}.$$

The linearisation of $\bar{\partial}_J$ at u is a map

$$(24) \quad D_{u,J} = D_u \bar{\partial}_J : \mathcal{T}_u^p \rightarrow \mathcal{V}_u^p,$$

where $\mathcal{T}_u^p = T_u \mathcal{B}^p(\nu)$ is the space of globally L_1^p sections of $u^* T^\vee \widehat{E}$ tangent to Q at the boundary.

Lemma 3.11 *When ν is non-degenerate, the map $D_{u,J}$ from (24) is Fredholm, with index independent of p .*

Sketch proof The Fredholm property in the quilted setting turns out to be no more difficult than that of the construction of Floer homology [10] and its continuation maps. We first discuss the situation when $p = 2$.

We can find a finite cover $\{U_\alpha\}_\alpha$ for \widehat{S}_Γ of the following type: each U_α is the image in \widehat{S}_Γ of a corresponding subset $U'_\alpha \subset \widehat{S}$. Moreover, $(U'_\alpha, U'_\alpha \cap \widehat{\Gamma}')$ is the image of an embedding of pairs whose domain is one of the following:

- (i) (D^2, \emptyset) ;
- (ii) $(S^1 \times [0, \infty), \emptyset)$ or $(S^1 \times (-\infty, 0], \emptyset)$;
- (iii) $(D^2, (-1, 1) \times \{0\})$;
- (iv) $(S^1 \times (-\infty, 0], \{e^{2\pi i k/m} : k = 1, \dots, m\} \times (-\infty, 0])$ or $(S^1 \times [0, \infty), \{e^{2\pi i k/m} : k = 1, \dots, m\} \times [0, \infty))$.

Let $i_\alpha : U_\alpha \rightarrow \widehat{S}_\Gamma$ be the inclusion. In case (iii), $i_\alpha^* \xi$ is a pair (ξ_+, ξ_-) of sections of bundles E_\pm over $D_\pm^2 = \{z \in D^2 : \pm \operatorname{im}(z) \geq 0\}$. But $\xi_+ + \xi_- \circ \tau$ then defines a section of $E_+ \oplus \tau^* E_- \rightarrow D_+^2$, where $\tau : D_+^2 \rightarrow D_-^2$ is complex conjugation. This section is subject to a Lagrangian boundary condition, the linearisation of Q . Standard elliptic theory then shows that for any precompact open set $U \subset \widehat{S}_\Gamma$ and any compact $K \subset U$, one has

$$(25) \quad \|i_\alpha^* \xi\|_{L_k^2(K)} \leq \operatorname{const.} (\|i_\alpha^* \xi\|_{L^2(K)} + \|i_\alpha^* D_{u,J} \xi\|_{L_{k-1}^2(U)}).$$

Now consider first an end of type (iv). One sees that in this case $i_\alpha^* \xi$ is a section of a bundle over the strip $\{e^{2\pi i t} : t \in [0, 1/m]\} \times (-\infty, 0]$, subject to non-degenerate linear Lagrangian boundary conditions. (The reduction is constructed by ‘pleating’, as in our treatment of almost complex structures.) We may take the bundle to be a trivial complex bundle, of rank r , say, and the boundary conditions to be two transverse totally real subspaces (Λ_0, Λ_1) . The operator $D_{u,J}$ reduces to one of shape $\frac{d}{ds} + L$, where $L = i \frac{d}{dt} : L_1^2([0, 1/m], \mathbb{C}^r; \Lambda_0, \Lambda_1) \rightarrow L^2([0, 1/m], \mathbb{C}^r)$ is formally self-adjoint and invertible. The operator $\frac{d}{ds} + L : L_1^2(\mathbb{R} \times [0, 1/m], \mathbb{C}^r; \Lambda_0, \Lambda_1) \rightarrow L^2(\mathbb{R} \times [0, 1/m], \mathbb{C}^r)$ is an isomorphism [7, Section 3.1]. The story for ends of type (ii) is similar. For both types, the L^2 -norm of $\xi \in \ker D_{u,J}$, over the infinite strip, is controlled by its L^2 -norm over a compact band. It follows that the unit ball in $\ker D_{u,J} \cap L^2$ is compact, hence that the L^2 kernel is finite-dimensional.

As to finite-dimensionality of the cokernel, the point is to construct a parametrix for $D_{u,J}$, i.e. a bounded linear operator $P : \mathcal{V}_u^2 \rightarrow \mathcal{T}_u^2$ such that $D_{u,J}P$ differs from the

identity map by a compact operator. The construction is done by patching parametrices P_α for sections supported in U_α . For the patching procedure (essentially that of [2, p. 54]) see [7, Section 3.2]. The existence of a parametrix in cases (i) and (iii) follows from elliptic theory. In cases (ii) and (iv) we take the inverse P_α to the invertible operator $d/ds + L$.

The L^p theory introduces some additional complications, but the hard work is in dealing with the strip itself [10, 7]. Elements of $\ker D_{u,J} \cap L^p$ decay exponentially over the ends, and so lie in any L^q . Similarly for the kernel of the formal adjoint operator. \square

We define an enlargement $\tilde{\mathcal{Z}}(\nu, J)$ of $\mathcal{Z}(\nu, J)$ as the space pairs $(u, [\gamma])$, where (i) $u \in \tilde{\mathcal{Z}}(\nu, J)$, (ii) γ is a finite disjoint union of trees of simple parametrised J -holomorphic spheres in fibres of π , rooted at points of $\text{im}(u)$. Thus $\tilde{\mathcal{Z}}(\nu, J)$ contains $\mathcal{Z}(\nu, J)$ as the subspace where γ is the empty union. The topology is that of Gromov convergence.

The pleating trick allows us to establish Gromov–Floer compactness, which says that any bounded-action sequence in $\tilde{\mathcal{Z}}(\nu)$ converges, in the Gromov–Floer topology, to a *broken section* plus finitely many trees of spherical bubbles in fibres plus finitely many trees of discs in boundary fibres. The point here is that the process of bubble formation is local (in the domain) and therefore the usual proofs need only cosmetic changes.

3.3.3 Definition of the cobordism-maps

Let O_- and O_+ be objects in \mathcal{C} , and \mathcal{F} a matched collection of fibrations defining a cobordism between them. Let J_-^v and J_+^v be regular complex structures defining the Floer complexes $\text{CF}_*(O_\pm; J_\pm^v)$. Also pick $J \in \mathcal{J}(\widehat{E}_\Gamma; \{J_\pm^v\})$. Given cycles of Hamiltonian chords ν_- and ν_+ for the respective objects, there is a locally constant action function A on $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J)$ with sub-level sets $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J)_{\leq c} := A^{-1}((-\infty, c])$. Assuming we are in the normalised monotone (or strongly negative) case, the zero-dimensional part $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J)_{\leq c}^0$ is compact, because its only possible ends arise from breaking trajectories, and that can only occur when the index is ≥ 1 . Hence the zero-dimensional part of the level set of c , $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J)_c^0 = A^{-1}(c) \cap \tilde{\mathcal{Z}}(\nu_-, \nu_+; J)^0$, is also compact. Put

$$(26) \quad n_J(\nu_- | \mathcal{F} | \nu_+) = \sum_{r \in \mathbb{R}} \# \tilde{\mathcal{Z}}(\nu_-, \nu_+; J)_r^0 t^r \in \Lambda_R.$$

These coefficients define a chain map $\text{CF}(\mathcal{F}): \text{CF}(O_-; J_-) \rightarrow \text{CF}(O_+; J_+)$:

$$(27) \quad \text{CF}(\mathcal{F})(\nu_-) = \sum_{\nu_+} n(\nu_- | \mathcal{F} | \nu_+) \langle \nu_+ \rangle.$$

These maps compose compatibly with sewing of fibrations. This is a consequence of the standard gluing principles for Cauchy–Riemann operators. We shall only invoke sewing along cylindrical ends, where the relevant gluing theorem is (in all analytical essentials) the one that underpins Hamiltonian Floer homology (cf. [26], for instance).

3.3.4 Quantum cap product

Assuming N monotone, the quantum cap product

$$H^p(N; \Lambda_R) \otimes_{\Lambda_R} \mathrm{HF}_*(\mu) \rightarrow \mathrm{HF}_{*-p}(\mu), \quad c \otimes x \mapsto c \frown x,$$

could in principle be defined by means of cocycles for any ordinary cohomology theory in which one can make sense of transversality. We opt for smooth singular cohomology.

A smooth singular q -chain σ in N is a formal linear combination $\sum_k a_k \sigma_k$, where $a_k \in \Lambda_R$ and $\sigma_k: \Delta^q \rightarrow N$ is a smooth q -simplex for each k . ‘Smooth’ means that σ_k extends to a C^∞ map defined on an open neighbourhood of Δ^q in \mathbb{R}^q , so in particular, the restriction of σ_k to a face of Δ^q is also smooth. These chains form a complex $S_*^{\mathrm{sm}}(N)$; its dual is the smooth singular cochain complex $S_{\mathrm{sm}}^*(N)$.

Take a partly ω_μ -compatible almost complex structure J on $\mathbb{R} \times T(\mu)$, eventually translation invariant and fully compatible over the two ends of the cylinder. For $\nu_\pm \in \mathcal{H}(T(\mu))$, one has the ‘continuation map’ moduli spaces $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J)$, and evaluation maps

$$\mathrm{ev}: \tilde{\mathcal{Z}}(\nu_-, \nu_+; J) \rightarrow N, \quad u \mapsto u(0, [0]).$$

Given any countable collection of smooth simplices $\sigma^k: \Delta^{d(k)} \rightarrow N$, we can find perturbations J' of J , supported over a neighbourhood of $(0, [0]) \in \mathbb{R} \times S^1$, so that the q -dimensional moduli spaces $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J')^q$ are transverse to the σ^k for all k and all $q \leq q_0$ (see [30, Lemma 2.5], for instance). One then has fibre products

$$(28) \quad \tilde{\mathcal{Z}}(\nu_-, \nu_+; J') \times_{\mathrm{ev}} \sigma^k.$$

Choose a codimension p smooth singular chain τ such that τ and $\partial\tau$ are transverse to the $\tilde{\mathcal{Z}}(\nu_-, \nu_+; J')^q$. By counting the points of fixed energies in the fibre products (28) one defines the matrix entries in a linear map

$$(29) \quad \tau \frown \cdot: \mathrm{CF}_*(\mu) \rightarrow \mathrm{CF}_{*-p}(\mu).$$

The usual trajectory-breaking argument demonstrates that

$$(30) \quad \partial(\tau \frown x) = (\partial\tau) \frown x + \tau \frown \partial x.$$

This argument requires transversality for the 1-dimensional fibre products, and the fact that these 1-dimensional spaces compactify to (topological) 1-manifolds with boundary. The latter fact is part of the gluing theory for holomorphic spheres [18]. To define the quantum cap product action of a smooth singular cochain c is defined to be that of its Poincaré dual $c \cap [N]$.

Remarks (i) In this paper we will only need the quantum cap-product action of a few specific chains. We have not defined a chain-level action of the *whole* of $S_*^{\text{sm}}(N)$.

(ii) There was nothing special about the base $\mathbb{R} \times S^1$ used in this construction; there is a similar story for any matched collection of fibrations with a marked point in the base (not on a seam), and it is compatible with sewing of these fibrations. We will not say any more about the formalities.

(iii) The quantum cap product also exists in the strongly negative case. The construction is easy, because if the almost complex structures are chosen generically one does not encounter bubbles at all (bubbles appear in generic 3-parameter families of almost complex structures).

4 Proof of the Gysin sequence: geometric aspects

4.1 Setting up the isomorphism

The maps $\rho^!$ and $e(V) \frown \cdot$ in the Gysin sequence (1) fit into the field theory described in the previous section. Let $\widehat{V}' = (\text{id}_{M_-} \times \mu)\widehat{V} \subset M_- \times N$. As promised in the introduction, we will construct an isomorphism

$$\Phi: H \text{ cone}(e(V) \frown \cdot) \rightarrow \text{HF}_*(\widehat{V}, \widehat{V}')$$

as the map on homology induced by a quasi-isomorphism

$$(31) \quad (H, C\rho): \text{cone}(e(V) \frown \cdot) \rightarrow \text{CF}_*(\widehat{V}, \widehat{V}').$$

So

$$C\rho: \text{CF}_*(\mu)[-k-1] \rightarrow \text{CF}_*(\widehat{V}, \widehat{V}')$$

is a chain map, and H a nullhomotopy of $C\rho \circ Ce$. Here Ce is the map defined by cap product with a cochain representing $e(V)$.

4.1.1 Global angular cochains

The Euler class of an oriented sphere bundle vanishes when pulled back to the total space:

$$(32) \quad \rho^* e(V) = 0.$$

When V is an $\mathrm{SO}(k+1)$ -bundle, this can be seen very simply: $\rho^* V$ has a tautological section. A more precise statement, valid for any smooth, R -oriented fibre bundle whose fibres are R -homology k -spheres, is that the orientation class o on the fibre F *transgresses* to the Euler class $e(V)$, where we work with coefficients in the ring R . Recall that the transgression τ is the homomorphism

$$d_{k+1}^{0,k} : E_{k+1}^{0,k} \rightarrow E_{k+1}^{k+1,0}$$

in the Leray–Serre spectral sequence $(E_r^{p,q}, d_r^{p,q})$ of a fibre bundle $\pi : E \rightarrow B$. One thinks of it as a map from a subgroup of $H^k(F; R)$ to a quotient of $H^{k+1}(B; R)$. The cochain interpretation is that $\tau([x]) = [y]$ when x extends to a cochain $w \in C^k(E; R)$ such that $dw = \pi^* y$.

In the case of an R -oriented homology- S^k -bundle $\rho : V \rightarrow N$, there is a transgression homomorphism $\tau : H^k(F; R) \rightarrow H^{k+1}(N; R)$. One has $\tau(o) = e(V)$ (indeed, $e(V)$ is sometimes defined this way). Hence there exists $\psi \in C^k(V; R)$ such that $d\psi = \rho^* e(V)$ and $[\psi|_{S^k}] = o$. We shall call such a cochain ψ a ‘global angular cochain’.⁴ Dually, if Z is a (codimension $k+1$) chain Poincaré dual to $e(V)$, one can find a ‘global angular chain’: a codimension k chain Y on V such that $\partial Y = \rho^! Z$ and $[Y \cap \rho^{-1}(x)] = [\text{pt.}]$.


References for this discussion are [33, 5]. A small technical concern is that Serre [33] sets up the spectral sequence via cubical singular (co)chains. However, from this one can readily deduce its analogue for the smooth singular (co)chains complexes used here.

4.1.2 Defining $C\rho$

The map $C\rho$ is associated with a matched collection of fibrations over a quilted surface depicted schematically in Figure 3. The base surface S is a finite cylinder; its elongation \widehat{S} (see (19)) an infinite cylinder. To be precise, we set $\widehat{S} = \mathbb{C}/4i\mathbb{Z}$, and

$$(33) \quad S = \{[z] \in \widehat{S} : |\operatorname{Re}(z)| \leq 2\},$$

⁴Bott and Tu call the ‘de Rham’ analogue of ψ a *global angular form* [5, p. 121].

Figure 3: The base of the matched pair of fibrations underlying $C\rho$.

then define $q: \mathbb{C} \rightarrow \widehat{S}$ as the quotient map. The arc $\widehat{\Gamma} \subset S$ is the union of two line segments and a circular arc⁵:

$$(34) \quad \widehat{\Gamma} = q(\{z : \operatorname{Im}(z) = \pm 1, \operatorname{Re}(z) \geq 1\}) \cup q(\{z : |z - 1| = 1, \operatorname{Re}(z) \leq 1\}).$$

Figure 3 shows the finite part $\Gamma := \widehat{\Gamma} \cap S$. The point p shown in the diagram (which is nothing more than a convenient reference point) is $q(0)$. The closures of the two components of $\widehat{S} \setminus \widehat{\Gamma}$ are denoted by B (the convex part) and C , as in the figure.

To construct an LHF over C , let $R = (N \times \mathbb{C})/\mathbb{Z}$ where $n \in \mathbb{Z}$ acts by $(x, z) \mapsto (\mu^n(x), z + 4in)$. There is a projection map $\pi_R: R \rightarrow \widehat{S}$, $[x, z] \mapsto [z]$, and a natural closed two-form Ω_R on R whose pullback to $N \times \mathbb{C}$ is $\operatorname{pr}_1^* \omega$. Now let $R = \pi_R^{-1}(C)$. Let

$$(35) \quad \pi_Q: Q \rightarrow C$$

be the restriction of π_R to Q , and Ω_Q the restriction of Ω_R to Q . Then (Q, π_Q, Ω_Q) is a flat LHF over C . Let

$$(36) \quad P = B \times M,$$

and let $\pi_P: P \rightarrow B$ be the trivial projection. Give P the closed 2-form

$$(37) \quad \Omega_P := \omega_M + d(Kdt)$$

for some function $K \in C^\infty(M)$. Here $z = s + it$ is the complex coordinate on \mathbb{C} (or on \widehat{S}). For later reference, we record the following easy observation.

Lemma 4.1 *(P, π_P, Ω_P) is a flat LHF over B .*

⁵Pedantically, we ought to smooth the joins between the line segments and the arc, but this detail will be suppressed.

The fibre $Q_p = \pi_Q^{-1}(p)$ is identified with N by thinking of it as the image of $N \times \{0\} \subset N \times \mathbb{C}$ under the \mathbb{Z} -action. The fibre $P_p = \pi_P^{-1}(p)$ is identified with M because π_P is by definition a trivial bundle. Our Lagrangian submanifold $\widehat{V} \subset M_- \times N$ appears here:

$$(38) \quad \widehat{V}_p := \widehat{V} \subset (P_p)_- \times Q_p = M_- \times N.$$

We denote this submanifold by \widehat{V}_p to emphasise that it lies in the product of the fibres over $p \in \widehat{S}$. Symplectic parallel transport along the arc $\widehat{\Gamma}$, starting at p , carries \widehat{V}_p to a Lagrangian $\widehat{V}_x \subset (P_x)_- \times Q_x$ for each $x \in \widehat{\Gamma}$. Write

$$(39) \quad \widehat{V} = \bigcup_{x \in \widehat{\Gamma}} \widehat{V}_x.$$

Definition 4.2 The map $C\rho$ is the cobordism-map (27) associated with the matched pair of fibrations (P, Q, \widehat{V}) . Thus it counts finite-action pairs (u, v) , where u is a section of π_P , v a section of π_Q , both pseudo-holomorphic for adapted almost complex structures (3.10), and $(u(x), v(x)) \in \widehat{V}_x$ for each $x \in \widehat{\Gamma}$.

This definition needs fleshing out. The incoming boundary of S (which is the one on the left) is a circle without marked points, and over it we have an LHF isomorphic to the mapping torus $T(\mu) \rightarrow S^1$. The outgoing boundary (on the right) is a circle with two marked points, $z_1 = [2 - i]$ and $z_2 = [2 + i]$. Now, let us identify the fibre of P over z_1 with M via symplectic parallel transport along Γ , starting at p . Similarly, identify the fibre of Q over z_2 by with N via symplectic parallel transport along Γ from p . The restriction of P to the directed boundary-arc $[z_1, z_2]$ is trivialised by parallel transport starting at z_1 . The restriction of Q to the directed boundary-arc $[z_2, z_1]$ is trivialised by parallel transport starting at z_2 . By using these trivialisations we see that the target for $C\rho$ is $\text{CF}_*((\text{id} \times \mu^{-1})(\widehat{V}), \widehat{V}) = \text{CF}_*(\widehat{V}, (\text{id} \times \mu)(\widehat{V}))$. Thus this matched pair of fibrations defines a map

$$(40) \quad C\rho: \text{CF}_*(\mu) \rightarrow \text{CF}_*(\widehat{V}, (\text{id} \times \mu)(\widehat{V})).$$

4.1.3 Defining Ce

Choose a smooth singular chain Z representing the Poincaré dual to the Euler class $e(V)$. Define the quantum cap product map

$$Ce = Z \frown \cdot: \text{CF}_*(\mu) \rightarrow \text{CF}_{*-k-1}(\mu).$$

Let

$$e: \text{HF}_*(\mu) \rightarrow \text{HF}_{*-k-1}(\mu)$$

be the induced map on homology.




Figure 4

4.1.4 The composite

The gluing (or sewing) theorem for fixed-point Floer homology expresses the composite $C\rho \circ Ce$ in terms of a sewed fibration, as in Figure 4. Namely, Ce is defined using the LHF (R, π_R, Ω_R) over \widehat{S} defined above; $C\rho$ using a matched pair of LHF's over \widehat{S} . We can glue these together, as in the figure, to obtain a new matched pair of LHF's over \widehat{S} ; in fact, this new pair are naturally identified with P and Q . The only new feature is that we have an additional marked point in the interior of C . For definiteness, let us say that this marked point is $p' := q(-1)$. In the diagram, we have indicated the marked points where incidence conditions are specified with solid black dots; marked points which are needed just for reference are shown as hollow dots.

The composite $C\rho \circ Ce$ counts pairs of pseudo-holomorphic sections (u, v) as above, but now v sends the marked point p' to a point of $Z \subset M = \pi_Q^{-1}(p')$.

4.1.5 The nullhomotopy

The nullhomotopy H promised in (31) will be constructed as the concatenation of two homotopies, H_1 and H_2 . The first homotopy, H_1 , arises by moving the marked point from p' to p , as in Figure 5.

The matrix entries of this homotopy are defined as counts of isolated triples $\{(t, u, v)\}$, where (u, v) is a pair of pseudo-holomorphic sections as before, $t \in [-1, 0]$, and $u(q(t)) \in \phi_t(Z)$, where $\phi_t: Q_p \rightarrow Q_t$ is (leftward) symplectic parallel transport along

the path $q([t, 0])$. Formally,

$$(41) \quad H_1\langle \nu \rangle = \sum_{x \in \widehat{V} \cap \widehat{V}'} \sum_{r \in \mathbb{R}} \# \left(\widetilde{\mathcal{Z}}(\nu, x; J_r^k \times [-1, 0]) \times_{\text{ev}, \phi} (Z \times [-1, 0]) \right) t^r \langle x \rangle$$

where $\text{ev}(z, t) = u(q(t))$ and $\phi(\cdot, t) = \phi_t \circ Z$. This definition requires J (which amounts to a pair of complex structures (J_P, J_Q) on the two fibrations P and Q) to be regular in the sense that the ev is transverse to ϕ .

We have to take a little care over compactness: whilst spherical bubbles in the fibres are tolerable, disc-bubbles attached to \widehat{V} are not. This is where the assumption MAS comes into play. When $(t; u, v)$ contributes to the moduli space, the pair (u, v) has index $k + 1$; consequently, when \widehat{V} is monotone, disc-bubbles in the compactified moduli space have Maslov index $\leq k + 1$, and consequently do not occur when $m_{\widehat{V}}^{\min} \geq k + 2$. In the strongly negative case, we can choose our almost complex structures so as to have the generic property that there are no disc-bubbles at all (note that this can be done using product complex structures on $M_- \times N$, because these suffice for a transversality of evaluation arguments along the lines of [30, Lemma 2.10]).

Corollary 4.3 *We have $\partial H_1 + H_1 \partial = C\rho \circ Ce + f$, where f is the homomorphism defined via pairs (u, v) with $v(1) \in Z$.*

Proof This is almost an instance of a standard principle, but its non-standard aspect is that we initially mark the *interior* point $p' \in C$, but then slide this to the *boundary* point p . We have to show that, despite this, the (formally) 1-dimensional component \mathcal{Z}_1 of the moduli space of triples $(t; u, v)$, where $v(t) \in \phi_t(Z)$, is a 1-manifold with boundary. This comes down to showing that the evaluation map $\text{ev}: [0, 1] \times \mathcal{Z}_1 \rightarrow N$, $(t, u, v) \mapsto v(t)$, can be made transverse to the cycle Z .

First consider an index k moduli space $\mathcal{Z}_k(x, y)$ of pairs (u, v) without point constraints. We have an evaluation map $\text{ev}_1: \mathcal{Z}_k(x, y) \rightarrow \widehat{V}_1$. We can choose the complex structures on P and Q so as to make ev_1 transverse to the codimension k cycle $\rho^!Z$ in \widehat{V} , by an argument explained by Seidel [30, Lemma 2.5], for example.

It follows that $\text{pr}_N \circ \text{ev}_1$ is transverse to Z . This projected evaluation map is the same as the evaluation map ev_1^v , where we evaluate only the second component of the pair (u, v) . Now, $\text{ev}_1^v: \mathcal{M}_k(x, y) \rightarrow N$ depends smoothly on t . Since transversality is an open property, it follows that for small $\epsilon > 0$ and $t \in [1 - 2\epsilon, 1]$, ev_1^v is also transverse to Z . Combining this with an unproblematic transversality argument for $[0, 1 - \epsilon]$ now gives the result. \square



Figure 5

The second homotopy, H_2 , arises via a global angular cochain. As in (4.1.1), let $Y = \psi \cap [V] \in S_{2n-2k}^{\text{sm}}(V)$ be a global angular chain. Thus $\partial Y = \rho^! Z$ and $[Y \cap \rho^{-1}(x)] = [\text{pt.}] \in H_0(V; R)$. We use Y to define a nullhomotopy of f . Define

$$(42) \quad H_2: \text{CF}(\mu) \rightarrow \text{CF}(\widehat{V}, (\text{id}_M \times \mu)\widehat{V}), \quad c \mapsto \sum_x n_1(\widetilde{c}|\text{gr}(\rho|_Y)|x) \langle x \rangle.$$

The coefficient $n_1(\widetilde{c}|\text{gr}(\rho|_Y)|x)$ counts the isolated points in a one-parameter moduli space, weighted by t^ϵ with ϵ the action. We have

$$(43) \quad \partial H_2 + H_2 \partial = f.$$

Indeed, the right-hand side corresponds to pairs of sections (u, v) where $u(p)$ lie in the image of $\rho^! Z$ (that is, in $\rho^{-1}(\text{im } Z)$); but this is equivalent to $v(p)$ passing through Z . Hence

$$(44) \quad \partial(H_1 + H_2) + (H_1 + H_2)\partial = C\rho \circ Ce,$$

i.e. $H := H_1 + H_2$ is a nullhomotopy of $C\rho \circ Ce$.

4.2 Lagrangian intersections

The chain complexes defined in the previous section are well-defined up to homotopy-equivalence, and the maps between them are well-defined up to homotopy. The resulting homology modules are independent of choices up to isomorphism (in fact, canonical isomorphism) and the maps between them are canonical. However, the proof of Theorem 1 appeals to specific choices, to be elaborated in this subsection and the next.

Our first task is to analyse the intersections of \widehat{V} with a particular Hamiltonian perturbation of $(\text{id}_{M_-} \times \mu)(\widehat{V})$. To begin, $\widehat{V} \cap (\text{id}_{M_-} \times \mu)\widehat{V}$ is the set of pairs $(x, \rho x)$ such that $x \in V$ and $\rho x \in \text{Fix}(\mu)$:

$$(45) \quad \text{pr}_1 \left(\widehat{V} \cap (\text{id}_{M_-} \times \mu)\widehat{V} \right) = \rho^{-1}(\text{Fix}(\mu)).$$

After a C^k -small Hamiltonian perturbation, μ has non-degenerate fixed points, according to the following well-known lemma which follows from [15, Theorem 3.1].

Lemma 4.4 *There is a dense set of functions $h \in C^\infty(N)$ such that $\phi_h \circ \mu$ has non-degenerate fixed points, where ϕ_h is the time-one Hamiltonian flow.*

We choose such an h and rename $\phi_h \circ \mu$ as μ (this does not affect Floer homology). The intersection $\widehat{V} \cap (\text{id}_{M_-} \times \mu)\widehat{V}$ in $M_- \times N$ is then clean and diffeomorphic to $\text{Fix}(\phi_h \circ \mu) \times F$, where $F \simeq S^k$ is the typical fibre of ρ .

To make the intersection transverse, we shall make a controlled Hamiltonian perturbation to id_{M_-} . By the coisotropic neighbourhood theorem, the symplectic structure in a neighbourhood of $V \subset M$ is determined by that of N . For a specific model, choose a connection α in the fibre bundle $\rho: V \rightarrow N$. It determines an embedding $i_\alpha: T^*_V(V) = (T^V V)^* \rightarrow T^*V$ as a sub-bundle complementary to $\text{Ann}(T^V V)$. On the space

$$(46) \quad W := i_\alpha(T^*_V V) \subset T^*V,$$

we consider the symplectic form

$$(47) \quad \sigma = i_\alpha^* d\lambda_V + R^* \rho^* \omega_N,$$

where $R: W \rightarrow V$ is the projection to V and λ_V the canonical one-form on T^*V . Fix a Riemannian metric on V , and let $W_\epsilon \subset W$ be the subset of vertical cotangent vectors of length $< \epsilon$.

For small ϵ , there is a symplectomorphism $\text{nd}(V) \rightarrow W^\epsilon$ extending the zero-section $V \rightarrow W_\epsilon$; we may therefore consider W_ϵ as a neighbourhood of V .

The space W_ϵ , with its projection $p = \rho \circ R: W_\epsilon \rightarrow N$, is an LHF. Its Hamiltonian curvature (7) is determined by the curvature of α . It is convenient to choose the connection form α to be flat over a neighbourhood $U = \bigcup_{\bar{x} \in \text{Fix}(\mu)} U_{\bar{x}}$ of the finite set $\text{Fix}(\mu)$. The Hamiltonian connection is then flat over U , hence (after shrinking U) symplectically trivial. For each $\bar{x} \in \text{Fix}(\mu)$, choose a symplectically trivialising projection

$$(48) \quad t_{\bar{x}}: p^{-1}U_{\bar{x}} \rightarrow (T^*F)_{<\epsilon},$$

so that the symplectic form on $p^{-1}U_{\bar{x}}$ is $p^* \omega_N + t_{\bar{x}}^* d\lambda_F$.

For each $\bar{x} \in \text{Fix}(\mu)$, fix a bump function $\eta_{\bar{x}} \in C^\infty(N)$, supported in $U_{\bar{x}}$. Choose a Morse function m on F , and let $\hat{m} \in C^\infty(T^*F)$ be the function $\hat{m} = m \circ s$, where $s: T^*F \rightarrow F$ is the projection. The Hamiltonian flow $\phi_{\hat{m}}$ with respect to the canonical form $-d\lambda_{\text{can}}$ is ‘vertical’, and moves the zero section to the graph of dm .

Now consider functions $H \in C^\infty(M)$ of shape $\chi \cdot H_m$, where χ is a cutoff function for the neighbourhood W_ϵ of V , and

$$(49) \quad H_m = \sum_{\bar{x} \in \text{Fix}(\mu)} (R^* \rho^* \eta_{\bar{x}}) \cdot (\hat{m} \circ t_{\bar{x}}).$$

Lemma 4.5 *After rescaling m , all intersection points*

$$(x, \rho x) \in \widehat{V} \cap (\phi_{\chi H_m} \times \mu) \widehat{V}$$

satisfy $(x, \rho(x)) \in \text{Fix}(\phi_{\chi H_m} \times \mu)$.

Proof Since \bar{x} is a non-degenerate fixed point of μ , there are Euclidean coordinates v on $U_{\bar{x}}$ centred on \bar{x} , and a constant $C > 0$, such that

$$\|\mu(v) - v\| \geq C\|v\|$$

when $\|v\|$ is sufficiently small. We may assume that this inequality is valid in $\text{supp}(\eta_{\bar{x}})$. On the other hand, by rescaling m by small positive reals, $\phi_{\chi H_m}$ can be made arbitrarily C^0 -close to the identity. In particular, we may arrange that $|\rho(y) - \rho \circ \phi_{\chi H_m}(y)| \leq C/4$ for all $y \in W_\epsilon$.

Take a point $(x, \rho(x)) = (\phi_{H_m}(y), \mu \rho(y)) \in \widehat{V} \cap (\phi_{\chi H_m} \times \mu) \widehat{V}$. Then $\rho(x) \in \text{supp}(\eta_{\bar{x}})$ for some $\bar{x} \in \text{Fix}(\mu)$, since if not, $H_m \equiv 0$ near x , so $x = y$ and $\rho(x) \in \text{Fix}(\mu)$, which is absurd. We have

$$0 = \|\mu \rho(y) - \rho(x)\| = \|\mu \rho(y) - \rho \phi_{\chi H_m}(y)\| \geq \frac{1}{2} C \|\rho(y) - \rho \phi_{\chi H_m}(y)\|,$$

so $\rho(x) = \rho \phi_{\chi H_m}(y) = \rho(y)$. Thus $\rho(x) \in \text{Fix}(\mu)$. Now, $\phi_{\chi H_m}$ preserves the fibre $\rho^{-1}(\rho(x))$, and (identifying the fibre with $(T^*F)_{<\epsilon}$ via $t_{\rho(\bar{x})}$) it acts as $\phi_{\hat{m}}$. So for $\phi_{\chi H_m}(y)$ to be a point of V , we must have $x \in \text{crit}(m)$ and $x = \phi_{\chi H_m}(y) = y$. Thus $(x, \rho(x)) \in \text{Fix}(\phi_{\chi H_m} \times \mu)$, as required. \square

The intersection points in the previous lemma are transverse because m was Morse.

The functions H_m are not Morse, and it will be useful to observe that result continues to hold for a convenient class of Morse functions on M . Start with a Morse function $l \in C^\infty(N)$ such that each $\bar{x} \in \text{Fix}(\mu)$ occurs as a local maximum of l . On W_ϵ , put

$$(50) \quad K'_{l,m} = H_m + \delta \sum_{c \in \text{crit}(l)} (R^* \rho^* l) \cdot (\hat{m} \circ t_c)$$

where t_c are local trivialising projections (like the projections $t_{\bar{x}}$ from (48)), and $\delta \ll 1$. Put

$$(51) \quad K_{l,m} = |\cdot|^2 + K'_{l,m}$$

where $|\cdot|^2$ denotes the norm-squared of cotangent vectors to F . The functions $K_{l,m}$ are Morse; the critical points lie on V , and project to the critical points of l on N .

Lemma 4.6 *Lemma 4.5 remains true when H_m is replaced by $K'_{l,m}$ or $K_{l,m}$.*

Proof The proof is similar but a little more complicated.

Given $\epsilon > 0$, we can find a constant C' so that if $\text{dist}(y, \mu(y)) < C'$ then y lies within ϵ of a fixed point of μ (an easy consequence of the compactness of N). Hence, when δ is small enough, the following holds: take any point $(x, \rho(x)) = (\phi_{K'_{l,m}}(y), \mu\rho(y)) \in \widehat{V} \cap (\phi_{K'_{l,m}} \times \mu)\widehat{V}$; then $\rho(x)$ lies in $\bigcup_{\bar{x} \in \text{Fix}(\mu)} \text{supp}(\eta_{\bar{x}})$.

Making δ still smaller, we may suppose that $|\rho(y) - \rho \circ \phi_{K'_{l,m}}(y)| \leq C/4$ for all $y \in W_\epsilon$. It is still true that $\phi_{K_{l,m}} - |\cdot|^2$ preserves the fibre $\rho^{-1}(\rho(x))$ when $\rho(x) \in \text{Fix}(\mu)$. By the same argument as before, one sees that for any intersection point $(x, \rho(x)) \in \widehat{V} \cap (\phi_{K'_{l,m}} \times \mu)\widehat{V}$ one has $\rho(x) \in \text{Fix}(\mu)$.

Moreover, the functions $K'_{l,m}$ and $|\cdot|^2$ Poisson-commute, which implies that

$$\phi_{K_{l,m}} = \phi_{|\cdot|^2} \circ \phi_{K'_{l,m}}.$$

But $\phi_{|\cdot|^2}$ preserves V , and hence for any intersection point $(x, \rho(x)) \in \widehat{V} \cap (\phi_{K_{l,m}} \times \mu)\widehat{V}$ one has $\rho(x) \in \text{Fix}(\mu)$ (because this is true for $K'_{l,m}$).

Finally, take an intersection point $(x, \rho(x)) \in \widehat{V} \cap (\phi_{K'_{l,m}} \times \mu)\widehat{V}$ (or in $\widehat{V} \cap (\phi_{K_{l,m}} \times \mu)\widehat{V}$) with $\rho(x) \in \text{Fix}(\mu)$. Then $\rho(x) = \rho(y)$ and $x \in \text{crit}(m)$ as before. \square

Proposition 4.7 *When m is small in C^0 and has precisely two critical points, ρ induces a two-to-one map $\beta: \widehat{V} \cap (\phi_{K_{l,m}} \times \mu)\widehat{V} \rightarrow \text{Fix}(\mu)$.*

Proof Immediate from the previous lemma. \square

Remark β gives a bijection between the generators of the complexes $\text{cone}(e)$ and $\text{CF}_*(\widehat{V}, (\phi_{K_{l,m}} \times \mu)\widehat{V})$.

4.2.1 The definition of K

Recall our standing hypothesis that F has a Morse function with precisely two critical points. Thus, by an observation of Reeb (see [20]), F is homeomorphic to S^k .⁶ We are in a position to stipulate that, from now on,

K should be one of the functions $\chi K_{l,m}$, where $K_{l,m}$ is from (51), χ is a cutoff function for the neighbourhood W_ϵ of V , and $m \in C^\infty(F)$ has precisely two critical points.

⁶Conversely, by results due to Smale—the h-cobordism theorem and the non-existence of exotic 5-spheres [34]—any F homeomorphic to S^k with $k \neq 4$ has such a Morse function.

4.3 Almost complex structures

4.3.1 On M

We continue to work on $W \cong T_v^*V$ (46). Notice that the connection α induces isomorphisms

$$(52) \quad \begin{aligned} T^*W &\cong R^*(T^*V) \oplus R^*(T_v^*V) \\ &\cong R^*(T_v^*V) \oplus (R^*(T_v^*V) \oplus R^*\rho^*(T^*N)). \end{aligned}$$

We shall consider almost complex structures I which have a block decomposition

$$(53) \quad I = \begin{bmatrix} 0 & -\text{id} & 0 \\ \text{id} & 0 & 0 \\ 0 & 0 & R^*\rho^*J_N \end{bmatrix}$$

with respect to the splitting (52), for some almost complex structure J_N on N compatible with ω_N . One has $\sigma(Iu, Iv) = \sigma(u, v)$ and $\sigma(v, Iv) > 0$ for $v \neq 0$, so there is an induced metric g_I making the splitting orthogonal.

Lemma 4.8 *The gradient ∇f of a function on W , with respect to g_I , is tangent to the zero-section $V \subset W$ at $x \in V$ if and only if $df(x)$ annihilates the first summand $R^*(T_v^*V)_x$. In the case of the functions H_m , $K'_{l,m}$ and $K_{l,m}$, this holds for all $x \in V$.*

Proof The first assertion is easily checked. That dH_m annihilates the ‘vertical’ summand $R^*(T_v^*V)$ follows from the Leibnitz rule and the facts that both $d(\hat{m} \circ t_{\bar{x}})$ and $R^*\rho^*d\eta_{\bar{x}}$ kill the relevant summand. The same argument applies to $dK'_{l,m}$. Since $d(|\cdot|^2)$ vanishes along the zero-section V , $dK_{l,m}$ too annihilates $R^*(T_v^*V)$ along V . \square

Notation 4.9 Write $\mathcal{J}(M; V)$ for the space of complex structures $I \in \mathcal{J}(M)$, compatible with ω_M , whose restriction to $\text{nd}(V)$ have the block decomposition (53) for some identification of $\text{nd}(V)$ with W_ϵ and some connection α .

4.3.2 On the matched pair of fibrations

Recall that $\pi_P: P \rightarrow B$ is a trivial M -bundle, whilst $\pi_Q: Q \rightarrow C$ is an N -bundle. We endowed Q with the closed 2-form Ω_Q which makes it a flat LHF with monodromy μ around the puncture. We gave $P = B \times M$ the closed 2-form $\Omega_P = \omega_M + d(Kdt)$, where t is the vertical (or imaginary) coordinate on B and $K \in C^\infty(M)$; in Lemma 4.1 we observed that this LHF is also flat. Remember also that we built the Lagrangian matching condition \hat{V} using the symplectic parallel transport defined by Ω_P and Ω_Q .

We now stipulate conditions on our almost complex structures.

- C has an incoming cylindrical end $e: (-\infty, 0] \times \mathbb{R}/\mathbb{Z} \rightarrow C$, $(s, t) \mapsto (s - 2, 4t)$. The pullback $e^*Q \rightarrow (-\infty, 0] \times \mathbb{R}/\mathbb{Z}$ is also cylindrical: it is $\text{pr}_2^*T_\mu$. We choose an adapted, regular almost complex structure $J_{Q,\infty}$ on e^*Q , so that the chain complex $\text{CF}_*(\mu; J)$ is defined.
- The quilted surface \widehat{S} also has an outgoing cylindrical end $f: [0, \infty) \times \mathbb{R}/\mathbb{Z}$, embedded by $f(s, t) = (s + 2, 4t)$, with seams $[0, \infty) \times \{1/4\}$ and $[0, \infty) \times \{-1/4\}$. The matched pair of fibrations, pulled back by f , is also translation-invariant.
We choose a regular pair of adapted almost complex structures $J_{P,+\infty}$ and $J_{Q,+\infty}$ on the two LHF's. We take $J_{P,+\infty}$ to be of a special form, namely, $J_{P,+\infty}|_{N \times \{(s,t)\}} = \phi_t^* J_M$ for an almost complex structure $J_M \in \mathcal{J}(M; V)$. Here ϕ_t is the time- t Hamiltonian flow corresponding to the function K .
- We draw J_Q from the space of fully compatible almost complex structures (3.10) on Q . For J_P on $P = B \times M$ we are still more prescriptive. Choose a compatible almost complex structure J_M on M . Define $J_P = J_P(J_M, K)$ to be the unique fully compatible almost complex structure on the bundle $B \times M \rightarrow B$ which restricts to the fibre $M \times \{(s, t)\}$ as $\phi_t^* J_M$.
- Let $\mathcal{J}(M, Q)$ denote the set of pairs (J_M, J_Q) where $J_M \in \mathcal{J}(M; V)$ and J_Q is an almost complex structure on Q fully compatible with π_Q .

4.3.3 Tweaking the almost complex structures

The definition of $\mathcal{J}(M, Q)$ suffers from a predictable defect: it is so stringent that $\mathcal{J}(M, Q)$ might not contain any regular almost complex structures. To get around this, we will take some $(J_M, J_Q) \in \mathcal{J}(M, Q)$ and tweak it slightly, without disturbing the features which are useful to us.

Consider the region $U_B = \{z \in B : \text{Re}(z) \leq 1, 2|z - 1| > 1\}$ (see Figure 6). Define $\tau: U_B \rightarrow \mathbb{C}$ by $\tau(z) = 1 + (\bar{z} - 1)^{-1}$. The quotient map $q: \mathbb{C} \rightarrow \widehat{S} = \mathbb{C}/\mathbb{Z}$ is injective on $\tau(U_B)$ and maps it to C ; let $U_C \subset C$ denote $q \circ \tau(U_B)$. We shall write τ also for the induced map $U_B \rightarrow U_C$, which is a diffeomorphism, mapping $U_B \cap \Gamma$ to $U_C \cap \Gamma$ by the identity map. Notice that the closure $\overline{U_C}$ wraps all the way around the cylinder: one has $q(1 + i\mathbb{R}) \subset \overline{U_C \cup U_B}$.

The product fibration $X := \tau^*(P|_{U_B}) \times_{U_C} (Q|_{U_C}) \rightarrow U_B$ is an LHS over U_C , with 2-form $(\Omega_P)_- \oplus \Omega_Q$. The pair $(J_M, J_Q) \in \mathcal{J}(M, Q)$ induces a fully compatible almost complex structure I in X . We consider the space $\text{Tw}(J_M, J_Q)$ of fully compatible




Figure 6: Schematic of S , pictured as a rectangle with two sides identified. It is divided along Γ into its two regions B (white, with its shaded sub-region U_B) and C (dark grey, with its lighter-shaded sub-region U_C).

almost complex structures I' in X such that the support $\text{supp}(I - I') \subset \overline{U_B}$ is contained in U_B .

We then write $\mathcal{Z}(J_M, J_Q, I')$ for the moduli space of pairs (u, v) where u is J_M -holomorphic except over U_B , v is J_Q -holomorphic except over U_C , and over U_B the pair (τ^*u, v) is I' -holomorphic; and $\tilde{\mathcal{Z}}(J_M, J_Q, I')$ for its enlargement to pairs with bubbles attached. We will call $\mathcal{Z}(J_M, J_Q, I')$ the moduli space of I' -tweaked pseudo-holomorphic pairs.

4.4 Low-energy pseudo-holomorphic sections

We consider pairs (u, v) where u is a J_P -holomorphic section of π_P , v a J_Q -holomorphic section of π_Q , and $(u(x), v(x)) \in \widehat{V}_x$ when x lies in $\widehat{\Gamma}$. We first identify a subset of the J_P -holomorphic sections of π_P . The latter is, by construction, a trivial fibration, and a J_P -holomorphic section of P is a map $u: B \rightarrow M$ satisfying

$$u_s + (\phi_t^* J_M)(u)u_t = 0.$$

If we put $\tilde{u}(s, t) = \phi_t \circ u(s, t)$ then this equation becomes $\tilde{u}_s + J_M(\tilde{u}_t - X(\tilde{u})) = 0$, i.e.,

$$(54) \quad \tilde{u}_s + J_M \tilde{u}_t - \nabla K(\tilde{u}) = 0.$$

A class of solutions are those satisfying

$$(55) \quad \tilde{u}_t = 0, \quad \tilde{u}_s = \nabla K(\tilde{u}).$$

Let us call these *gradient sections*. From equations (54, 37), the action of a section u equals

$$\begin{aligned}
 A(u) &= \int_B u^* \Omega_M = \int_B \omega_M(u_s, u_t - X(u)) ds dt \\
 &= \int_B \omega_M(\tilde{u}_s, \tilde{u}_t - X(\tilde{u})) ds dt \\
 &= \int_B |\tilde{u}_s|^2 ds dt,
 \end{aligned}
 \tag{56}$$

so for a gradient section one has

$$A(u) = \int_B |(\nabla K)(\tilde{u})|^2 ds dt.
 \tag{57}$$

In particular, if $e(K) = \inf_u A(u)$, where u ranges over gradient sections, then $e(K) < \infty$ and $e(\lambda K) = \lambda^2 e(K)$ for $\lambda \in \mathbb{R}$.

We require that $u(x) \in V_x$ when $x \in \hat{\Gamma}$; for a gradient section to satisfy this condition we need $(\nabla K) \circ (\tilde{u})$ to be tangent to V . But from Lemma 4.8, when $J_M \in \mathcal{J}(M; V)$, ∇K is *everywhere* tangent to V .

We can precisely identify the moduli space $\mathcal{M}^{\text{grad}}(J_M, J_Q)$ of pairs $(u, v) \in \mathcal{Z}(J_M, J_Q)$, where u is a gradient section and v is *horizontal*: there is a canonical identification

$$\mathcal{M}^{\text{grad}}(J_M, J_Q) \cong \bigcup_{x \in \text{crit}(K|_V): \rho(x) \in \text{Fix}(\mu)} W^u(K|_V; x),
 \tag{58}$$

where $W^u(K|_V; x)$ is the unstable manifold of x as a critical point of $K|_V$. This identification sends $y \in W^u(K|_V; x)$ to the unique pair $(u, v) \in \mathcal{M}^{\text{grad}}(J)$ such that $u(p) = y$ and v is the horizontal section corresponding to $\rho(x) \in \text{Fix}(\mu)$.

The unstable manifolds $W^u(K|_V; x)$ lie inside the fibres of ρ , and coincide with the unstable manifolds $W^u(m; x)$ of the Morse function m on the fibre. Thus

$$\mathcal{M}^{\text{grad}}(J) \cong \text{Fix}(\mu) \times \bigcup_{x \in \text{crit}(m)} W^u(m; x),$$

and in particular, $\mathcal{M}^{\text{grad}}(J)$ is stratified by the Morse indices of critical points of m .

The local dimension d of $\mathcal{M}^{\text{grad}}(J)$ is equal to $k - i(x)$, where $i(x)$ the Morse index of $x \in \text{crit}(m)$. The d -dimensional part $\mathcal{M}^{\text{grad}}(J)$ is cut out transversely, and its Zariski tangent space $T_{(u,v)} \mathcal{M}^{\text{grad}}(J)$ at a pair (u, v) is contained in $T_{(u,v)} \mathcal{M}(J)$, the Zariski tangent space of $\mathcal{M}(J)$.

Lemma 4.10 *For all $(u, v) \in \mathcal{M}^{\text{grad}}(J)$, one has $\ker D_{(u,v)} = T_{(u,v)} \mathcal{M}^{\text{grad}}(J)$ and $\text{coker } D_{(u,v)} = 0$. Hence $\text{ind } D_{(u,v)} = k - i(x)$.*

Proof This follows from the Weitzenböck formulae for Cauchy–Riemann operators in a Hermitian vector bundle. Working in $P \rightarrow B$, with the symplectic form $\Omega_P + \pi_P^*(ds \wedge dt)$, and considering the natural linearised symplectic connection in ∇^u in u^*T^vP , we have

$$(59) \quad D_u^* D_u X = \frac{1}{2} \nabla^{u*} \nabla^u X + \frac{J}{2} F_{\nabla}(\partial_s, \partial_t) X, \quad X \in C_c^\infty(B; u^*T^vP)$$

$$(60) \quad D_u D_u^* \xi = \frac{1}{2} \nabla^u (\nabla^u)^* \xi + \frac{J}{2} F_{\nabla}(\partial_s, \partial_t) \xi, \quad \xi \in \Omega_c^{0,1}(B; u^*T^vP)$$

by a computation that can conveniently be done using the formulae of [18, App. C]. The curvature term $F_{\nabla}(\partial_s, \partial_t)$, a section of the endomorphism bundle $\mathfrak{sp}(u^*T^vP)$, was identified in (9); it vanishes identically when F_Ω does. Similar formulae apply to D_v and its formal adjoint. Since both our LHF's are flat, the curvature terms in the Weitzenböck formulae vanish. Now, $\ker D_{(u,v)}$ is the subspace of $\ker D_u \times \ker D_v$ on which the matching condition holds. Whilst (59, 60) initially apply only to compactly supported sections, it follows by continuity that they are valid for smooth, exponentially decaying sections, in particular to elements of $\ker D_u$. Thus we see $T_{(u,v)} \mathcal{M}^{\text{grad}}(J) \subset \ker D_{(u,v)}$. The reverse inclusion is automatic, hence equality holds. Similarly, to show that $\text{coker } D_{(u,v)} = 0$, it suffices to show that $\ker D_{(u,v)}^* = 0$. But $\ker D_{(u,v)}^* = \ker D_u^* \times \ker D_v^*$, which by the formulae is equal to $\ker \nabla_u^* \times \ker \nabla_v^*$. But $\ker \nabla_v^* = 0$ because K is Morse–Smale, and $\ker \nabla_u^* = 0$ because μ is non-degenerate. \square

Lemma 4.11 *The number $\#\mathcal{M}^{\text{grad}}(\bar{x}, x)_0$ of isolated gradient-type pairs asymptotic to $\bar{x} \in \text{Fix}(\mu)$ and to $x \in \widehat{V} \cap \widehat{V}'$ is equal to 1 if $\rho(x) = \bar{x}$ and x is a maximum for m , and zero otherwise.*

Proof This is an immediate consequence of the structure of $\widehat{V} \cap \widehat{V}'$ (Proposition 4.7) together with the identification (58). \square

Now we consider the gradient-type trajectories in the moduli spaces defining the homotopy H . The moduli spaces for the first homotopy, H_1 , involve pairs (u, v) where v hits a cycle Z_t in some fibre Q_t , where t moves from p' to p . For generically-chosen cycles, no point in these moduli spaces will be a gradient-type pair (for this amounts to a finite-dimensional intersection problem in M , of a codimension k submanifold intersecting a 0-manifold). The second homotopy, H_2 , is more interesting. The contributions of gradient-type pairs to H_2 come from fibre products $\mathcal{M}^{\text{grad}} \times_{\text{ev}_p} Y$ where Y is the global angular chain bounding Z . One has

$$\mathcal{M}^{\text{grad}}(\bar{x}, x) \times_{\text{ev}_p} Y \cong W^u(x; m) \cap Y.$$

By construction, the algebraic intersection number $\rho^{-1}(\bar{x}) \cap Y$ is 1 for each $\bar{x} \in N$. So as to avoid issues concerning orientations of moduli spaces, we note that we can choose Y so that $\rho^{-1}(\bar{x}) \cap Y$ is transverse intersection in precisely one point whenever $\bar{x} \in \text{Fix}(\mu)$. We can then make the following conclusion.

Lemma 4.12 $\#\mathcal{M}^{\text{grad}}(\bar{x}, x) \times_{\text{ev}_p} Y = 1$ when $\rho(x) = \bar{x}$ and x is a minimum for m , while this number is zero in all other cases.

Lemma 4.13 Fix $(J_M, J_Q) \in \mathcal{J}(M, Q)$. Then we can find a C^∞ -dense set of $I' \in \text{Tw}(J_M, J_Q)$ so that $\mathcal{Z}(J_M, J_Q, I') \setminus \mathcal{M}^{\text{grad}}(J_M, J_Q)$ is cut out transversely.

Proof The LHF's (P, π_P, Ω_P) and (Q, π_Q, Ω_Q) have the simplifying feature that they are flat. Notice also that the set $U' := \overline{U_B} \cup \overline{U_C} \subset \widehat{S}$ has the property that any pair of horizontal sections (u, v) over U' (or rather, over its image in \widehat{S}_Γ), satisfying the matching condition \widehat{V} , extends to an element of $\mathcal{M}^{\text{grad}}(J_M, J_Q)$. By the argument of [30, Lemma 2.26], any horizontal (J_M, J_Q) can be tweaked by an $I' \in \text{Tw}(J_M, J_Q)$ so that all non-horizontal pairs in $\mathcal{Z}(J_M, J_Q, I')$ (or indeed in $\widetilde{\mathcal{Z}}(J_M, J_Q, I')$) are regular. \square

We now argue that it can be arranged that the contributions of the gradient-type pairs account for all the pseudo-holomorphic pairs of sufficiently low action. We need to make the notation more explicit, writing not just J_P but $J_{P,K}$ for the almost complex structure on P associated with the function K , and $\widehat{V}(K)$ for the matching condition associated with K .

Lemma 4.14 Choose a function K as in (4.2.1). There exists a constant $\epsilon > 0$ such that the following holds. Replace K by λK for some $\lambda \in (0, \epsilon]$. Choose $I' \in \text{Tw}(J_{P,\lambda K}, J_Q)$ and suppose that I' is a perturbation of size $\leq \epsilon$ in Floer's C^∞_ϵ -norm. Then, if (u, v) is any pair of finite-action, regular, I' -tweaked (J_P, J_Q) -holomorphic sections, subject to the matching condition $\widehat{V}(\lambda K)$, with $A(u) + A(v) \leq \epsilon$, we have $(u, v) \in \mathcal{M}^{\text{grad}}(J_M, J_{P,\lambda K})$.

Proof Suppose the lemma is not true. Let $J_{P,n}$ be the version of J_P constructed using $\frac{1}{n}K$ in place of K . Then there is a sequence of tweakings I'_n , perturbing $(J_{P,n}, J_Q)$ by an amount which goes to zero in C^∞_ϵ as $n \rightarrow \infty$, and of regular, I'_n -tweaked $(J_{P,n}, J_Q)$ -holomorphic pairs (u_n, v_n) , such that (i) $A(u_n) + A(v_n) \rightarrow 0$, and (ii) neither u_n nor v_n is horizontal. It is better to replace u_n by \tilde{u}_n , so that Floer's equation (54) holds, because (\tilde{u}_n, v_n) is subject to a Lagrangian matching condition which is

independent of n . Moreover, and satisfies a version of Floer's equation (54) whose coefficients converge in C^∞ as $n \rightarrow \infty$. Gromov–Floer compactness implies that some subsequence of (\tilde{u}_n, v_n) converges to a pair $(\tilde{u}_\infty, v_\infty)$ such that $A(u_\infty) = A(v_\infty) = 0$. Because the limiting energies are zero, no bubbling or trajectory-breaking is possible. Thus $u_\infty : B \rightarrow M$ must be a constant map, and v_∞ must be horizontal.

Up to this point of the paper, we have worked exclusively with non-degenerate matched fibrations. Now we need briefly to invoke Morse–Bott methods so as to handle the case $\lambda = 0$. It is well-known that the linear Fredholm theory underpinning Floer homology extends to the Morse–Bott case, provided one introduces exponentially weighted Sobolev spaces $L^p_{k,\delta}$ (see [7], for instance). This more general theory extends without difficulty to our setting of matched fibrations. We take $p = 2$ (so we have a Hilbert space) and take the weights δ very small, with signs chosen so that $L^2_{k,\delta}$ -sections may increase at a mild exponential rate over the ends. Thus, if D_0 denotes the deformation operator at (u_∞, v_∞) , where $\lambda = 0$, then $\ker D_0 \cong \mathbb{R}^k$ is the space of pairs which are zero over C and constant over B .

We consider the 1-parametric moduli space \mathcal{Z}^{par} of triples (λ, u, v) , where $\lambda \in \mathbb{R}$ and (u, v) is an index 0 holomorphic pair for the function λK . It is the zero set of a non-linear map Φ , which takes the following shape:

$$\Phi(\lambda, w) = \frac{1}{2} (dw + \alpha_\lambda \circ J \circ \alpha_\lambda^{-1} \circ dw \circ j), \quad \lambda \in \mathbb{R}.$$

Here we consider (u, v) as a section w of $E_\Gamma \rightarrow S_\Gamma$, while $\{\alpha_t\}$ is a 1-parameter family of automorphisms of TE_Γ starting at $\alpha_0 = \text{id}$. Moreover, α_t acts as the identity on $T^h E_\Gamma$. The deformation operator $\mathcal{D} = D_{(0,w)}\Phi$ is given by

$$\mathcal{D}(t, \dot{w}) = D_0 \dot{w} + \frac{t}{2} (\beta \circ J - J \circ \beta) \circ dw \circ j,$$

where $\beta = (d/dt)(\alpha_t)|_{t=0}$. But $\beta(T^h E_\Gamma) = 0$, and so if $w = (u_\infty, v_\infty)$ is our pair of horizontal sections then $(\beta \circ J - J \circ \beta) \circ dw = 0$. Hence $\ker \mathcal{D} = \ker D_0 \oplus \mathbb{R}$. We have $(\text{im } \mathcal{D})^\perp = (\text{im } D_0)^\perp \cong \mathbb{R}^k$.

A neighbourhood of $(0, u_\infty, v_\infty)$ in \mathcal{Z}^{par} can be described using a Kuranishi model. The neighbourhood is homeomorphic to a small neighbourhood of 0 in $\kappa^{-1}(0)$, where κ is the Kuranishi map

$$\kappa : \ker \mathcal{D} \rightarrow (\text{im } \mathcal{D})^\perp, \quad (t, \dot{w}) \mapsto \Pi \circ \Phi(t, \exp_{g_t} \dot{w}).$$

Here Π is orthogonal projection to $(\text{im } \mathcal{D})^\perp$ and g_t is a 1-parameter family of metrics. We have

$$\kappa(t, \dot{w}) = \Pi \circ \frac{t}{2} (\beta J - J \beta) \circ \dot{w} \circ j + O(|\dot{w}|^2 + t^2).$$

Write A for the linear map $\ker D_0 \rightarrow \operatorname{im}(\mathcal{D})^\perp$, $\dot{w} \mapsto \Pi \circ \frac{1}{2}(\beta J - J\beta) \circ \dot{w} \circ j$. Using the inverse function theorem, we can find coordinates charts near the origins of $\ker D_0$ and of $\operatorname{im}(\mathcal{D})^\perp$ so that $\kappa: \mathbb{R} \times \ker D_0 \rightarrow \operatorname{im}(\mathcal{D})^\perp$ takes the form

$$\kappa(t, x) \mapsto tA(x).$$

We argue that A must be a linear isomorphism. For it maps between spaces of the same dimension, and if it had a kernel, the fibres $\mathcal{Z}_\lambda^{\text{par}}$ of the projection $\mathcal{Z}^{\text{par}} \rightarrow \mathbb{R}$ would have points of positive local dimension for small $\lambda \neq 0$. Moreover, for $n \gg 0$, (n^{-1}, u_n, v_n) would be such a point. This contradicts regularity of (u_n, v_n) .

Thus A is an isomorphism, and there exist coordinates in which κ is given by the map

$$\mathbb{R} \times \mathbb{R}^k \rightarrow \mathbb{R}^k, \quad (\lambda, x) \mapsto \lambda x,$$

and $\text{pr}_2: \mathbb{R} \times \mathbb{R}^k \rightarrow \mathbb{R}$ corresponds to the projection $\mathcal{Z}^{\text{par}} \rightarrow \mathbb{R}$. The strand $\mathbb{R} \times \{0\} \subset \kappa^{-1}(0)$ parametrises gradient sections. For large n , (u_n, v_n) lies in this strand, which is what we want. \square

5 Proof of the Gysin sequence: algebraic aspects

5.1 \mathbb{R} -graded homological algebra

The algebraic mechanism we use to prove Theorem 2 is closely related to that used by Seidel to establish the exactness of the sequence describing the effect of Dehn twists on Floer homology. However, because the the symplectic action functional is not exact and the action spectrum not necessarily discrete, two new ingredients are needed: a completion (equivalent to linearity over the Novikov ring Λ_R), and the observation that the low-order terms are defined over the base ring R .

If R is a commutative unital ring, an \mathbb{R} -graded R -module is a R -module V with a given direct sum decomposition $V = \bigoplus_{r \in \mathbb{R}} V_r$. Its *support* is $\operatorname{supp}(V) := \{r \in \mathbb{R} : V_r \neq 0\}$. Given an interval $I \subset \mathbb{R}$, we say that V has *gap* I if $r, s \in \operatorname{supp}(V)$ implies $|r - s| \notin I$. A homomorphism $f: V \rightarrow V'$ between \mathbb{R} -graded modules has *order* I if $f(V_r) \subset \bigoplus_{s \in I} V'_{r+s}$ for all r .

The following lemma [30, Lemma 2.31] is Seidel's variation on a well-known principle.

Lemma 5.1 *Suppose that (D, δ) is a finitely-supported \mathbb{R} -graded R -module with gap $[\epsilon, 2\epsilon)$ for some $\epsilon > 0$, equipped with a differential of order $[0, \infty)$; thus we may write $\delta = \delta_{\text{low}} + \delta_{\text{high}}$, where δ_{low} is a differential of order $[0, \epsilon)$, and δ_{high} a map of order $[2\epsilon, \infty)$. Then, if δ_{low} is acyclic, so is δ ; that is, $H(D, \delta_{\text{low}}) = 0$ implies $H(D, \delta) = 0$.*

One can weaken the finite-support condition by demanding instead that $\text{supp}(D)$ is bounded above, but the condition cannot be dropped altogether: let $D = R[x] \oplus R[y]$. Let D_r be zero except when $r \in \mathbb{Z}_{\geq 0}$, in which case it is spanned by x^r and y^r . Let $\delta(f(x), g(y)) = (0, (1 - y)f(y))$. Then (taking $0 < \epsilon < 1$) δ_{low} is acyclic, but δ is not: the homology $H(D, \delta)$ is one-dimensional, spanned by $(0, 1)$.

The *completion* $\hat{V} = \varprojlim_r V/V_{\geq r}$ of the \mathbb{R} -graded module V is the submodule of $\prod_r V_r$ consisting of functions $v: \mathbb{R} \rightarrow V$, with $v(r) \in V_r$, such that $\text{supp}(v)$ has the ‘Novikov’ property that $\text{supp}(V) \cap (-\infty, c]$ is finite for any $c \in \mathbb{R}$. Completion is functorial: a homomorphism f of degree $[0, \infty)$ extends ‘by continuity’ to a homomorphism \hat{f} between completions. We can still speak of a map $F: \hat{V} \rightarrow \hat{V}'$ having order I : we mean that, for all r , we have

$$F(\hat{V}_r) \subset \left(\bigoplus_{s \in I} V'_{r+s} \right)^{\wedge},$$

where $\hat{V}_r \subset \hat{V}$ is the image of V_r under $V \rightarrow \hat{V}$, and the hat on the right-hand side denotes completion inside \hat{V}' .

What remains of the lemma when one drops the finiteness (or boundedness from above) hypothesis on $\text{supp}(D)$, but retains the gap assumption, is that $H(D, \delta_{\text{low}}) = 0$ implies $H(\hat{D}, \hat{\delta}) = 0$. (In the example with $D = k[x] \oplus k[y]$, one has $\hat{D} = k[[x]] \oplus k[[y]]$ and $(0, 1) = \hat{\delta}(0, \sum_{n \geq 0} y^n)$.) This assertion readily follows from Seidel’s lemma and the following observation.

Lemma 5.2 *Let (D, δ) be a differential R -module, $\{F^q D\}$ a decreasing filtration. Suppose that $H(D/F^q D) = 0$ for all q . Then $H(\varprojlim_q D/F^q D) = 0$.*

The proof is a simple exercise. (An unnecessarily fancy proof is to observe that both $\{D/F^q D\}$ and $\{H(D/F^q D)\}$ form inverse systems satisfying the Mittag–Leffler condition, so $\varprojlim^1 H(D/F^q D) = \varprojlim H(D/F^q D) = 0$, and hence $H(\varprojlim D/F^q D) = 0$, cf. [36].)

If one weakens the gap assumption, precautions are needed to safeguard even the conclusion that the completed \mathbb{R} -graded module is acyclic.

For A an \mathbb{R} -graded module, let \underline{A} denote the direct sum of shifts $\bigoplus_{r \in \mathbb{R}} A[r]$. If A has a differential d then \underline{A} then has an induced differential \underline{d} which respects the direct sum and restricts to $A[r]$ as d .

Lemma 5.3 *Suppose that (A, d) is a finitely-supported \mathbb{R} -graded module with gap $[\epsilon, 2\epsilon)$, equipped with a differential of order $[0, \epsilon)$, for some $\epsilon > 0$. Let $D = \underline{A}$. Let δ be a differential on \hat{D} such that*

- (i) δ is invariant under the \mathbb{R} -action on \hat{D} by shifts and continuous with respect to the completion.
- (ii) δ maps $A \subset \hat{D}$ into $\widehat{\bigoplus_{r \geq 0} A[-r]}$.
- (iii) $\delta = \delta_{\text{low}} + \delta_{\text{high}}$, where δ_{low} is equal to \widehat{d} , the differential induced by d , and δ_{high} is a homomorphism of order $[2\epsilon, \infty)$.
- (iv) $H(A, d) = 0$.

Then $H(\hat{D}, \delta) = 0$.

Remark We can think of \underline{A} as the tensor product $A \otimes_R R[\mathbb{R}]$, so that elements of $A[-r]$ can be written as $a \otimes t^r$, $a \in A$. When A has bounded support, the completion $\widehat{\underline{A}}$ of \underline{A} is naturally identified with $A \otimes_R \Lambda_R$. Then (i) just means that Λ_R -linear, and \widehat{d} is the Λ_R -linear extension of d .

Proof Assumption (iv) implies that $H(\underline{A}, \underline{d}) = 0$. It then follows from (iii) that $H(D, \delta_{\text{low}}) = 0$. Set

$$F^p D = \widehat{\bigoplus_{r \geq p\epsilon} A[-r]}$$

Thus $F^p D \supset F^{p+1} D$ and, by (i) and (ii), $\delta(F^p D) \subset F^p D$. This decreasing filtration of D induces a singly-graded spectral sequence $(E_k^*, d_k)_{k \geq 0}$ with $E_1^* = H(\text{gr}_*(D))$ (where gr_* denotes the associated graded module $\bigoplus_p F^p D / F^{p+1} D$). It follows from (iii) that $E_1^* = 0$. Indeed, A is the direct sum of ‘clumps’ C^i , preserved by d , whose support has width $< \epsilon$, and $F^p D / F^{p+1} D$ has corresponding clumps $\bigoplus_{r \in [p\epsilon, (p+1)\epsilon)} C^i[-r] \subset \bigoplus_{r \in [p\epsilon, (p+1)\epsilon)} A = F^p D / F^{p+1} D$. By (iii), the differential on the clumps is induced by the differential δ_{low} on C_i , which is acyclic. A spectral sequence argument as in [30, Lemma 2.31] then shows that $F^p D / F^{p+1} D$ is acyclic.

The spectral sequence E_k^* might not converge. However, if we look at the quotient $D / F^q D$, with its bounded-below, exhaustive filtration induced by the filtration of D , we obtain a spectral sequence converging to $H(D / F^q D)$. The E_1^* page is still zero. Thus $H(D / F^q D) = 0$. By the previous lemma, $H(\varprojlim_q D / F^q D) = 0$, i.e., $H(\hat{D}) = 0$. \square

Let us give property (ii) a name: given R -modules A and B and a Λ_R -linear map $f: A \otimes \Lambda_R \rightarrow B \otimes \Lambda_R$, let us say f is *positive* if $f(a) = \sum b_i t^{r_i}$ for elements $b_i \in B$ and $r_i \geq 0$.

The following lemma is a generalisation of a rotated version of [30, Lemma 2.32].

Lemma 5.4 (Double mapping cone lemma.) *Fix $\epsilon > 0$. Suppose that A, A' and A'' are finitely supported \mathbb{R} -graded R -modules satisfying*

GAP: A' and A have gap $(0, 3\epsilon)$ while A'' has gap $(0, 2\epsilon)$. If $r \in \text{supp}(A')$ and $s \in \text{supp}(A)$ then $|s - r| \notin (0, 4\epsilon)$.

Let $C = A \otimes_R \Lambda_R$, $C' = A' \otimes_R \Lambda_R$ and $C'' = A'' \otimes_R \Lambda_R$, and suppose that these three \mathbb{R} -modules are equipped with Λ_R -linear differentials $d_C, d_{C'}$ and $d_{C''}$. Suppose that

$$C' \xrightarrow{b} C \xrightarrow{c} C''$$

is a sequence of differential maps, and $h: C' \rightarrow C''$ a null-homotopy of $c \circ b$. Assume

POS: The maps $d_C, d_{C'}, d_{C''}, b, c$ and h are all positive.

ORD 1: The differentials $d_C, d_{C'}$ and $d_{C''}$ each have order $[2\epsilon, \infty)$. The map b has order $[2\epsilon, \infty)$; c and h have order $[0, \infty)$.

ORD 2: One can write $c = c_{\text{low}} + c_{\text{high}}$ and $h = h_{\text{low}} + h_{\text{high}}$, where the linear maps c_{low} and h_{low} have order $[0, \epsilon)$ while c_{high} and h_{high} have order $[2\epsilon, \infty)$.

Further assume

LOW 1: the maps c_{low} and h_{low} are induced by maps $c_0: A \rightarrow A''$ and $h_0: A' \rightarrow A''$, respectively; and

LOW 2: the map $(h_0, c_0): A \oplus A' \rightarrow A''$ is a linear isomorphism.

The induced differential map

$$(h, c): \text{cone}(b) \rightarrow C'',$$

is certainly positive and has order $[0, \infty)$ when we consider the mapping cone as the completion of the direct sum $C' \oplus C$ of \mathbb{R} -graded modules. The conclusion of the lemma is that (h, c) is a quasi-isomorphism.

Proof We shall invoke the last lemma to prove that the ‘double mapping cone’ $\text{cone}(h, c)$ is acyclic. As an \mathbb{R} -module, $\text{cone}(h, c) = C' \oplus C \oplus C'' = (A' \oplus A \oplus A'') \otimes \Lambda_R$, so requirement (i) of the previous lemma holds. Its differential δ has the block form

$$\delta = \begin{bmatrix} d_{C'} & 0 & 0 \\ b & -d_C & 0 \\ -h & c & d_{C''} \end{bmatrix},$$

and is positive by POS, so (ii) holds. We use ORD 2 to decompose the differential as $\delta = \delta_{\text{low}} + \delta_{\text{high}}$, where

$$\delta_{\text{low}} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -h_{\text{low}} & c_{\text{low}} & 0 \end{bmatrix}, \quad \delta_{\text{high}} = \begin{bmatrix} d_{C'} & 0 & 0 \\ b & -d_C & 0 \\ -h_{\text{high}} & c_{\text{high}} & d_{C''} \end{bmatrix}.$$

By ORD 1 and ORD 2, δ_{low} has order $[0, \epsilon)$, and squares to zero, while δ_{high} has order $[2\epsilon, \infty)$. By LOW 1, δ_{low} is induced by a differential d on $A' \oplus A \oplus A''$. The GAP conditions imply that $A' \oplus A \oplus A''$ has gap $[\epsilon, 2\epsilon)$. This deals with (iii). Finally, (iv) holds because of assumption LOW 2, and the proof is complete. \square

5.2 Completing the proof

All the ingredients in our proof of Theorem 1 are now to close to hand. Recall from the introduction that the Gysin sequence is derived from the exact sequence of a mapping cone: indeed, it is an immediate consequence of Theorem 2. The definition of ρ_* is algebraic. To prove Theorem 2 we shall invoke the double mapping cone lemma (5.4), taking the following as our input data.

- $A = A' = R\text{Fix}(\phi_h^{-1} \circ \mu)$, with \mathbb{R} -gradings concentrated in degree zero.
- $A'' = R(\widehat{V} \cap \widehat{V}')$, with \mathbb{R} -gradings still to be specified.
- The differential $d_C = d_{C'}$ on $C = C' = A \otimes_R \Lambda_R = \text{CF}(\mu)$ is the Floer-theoretic differential, for some regular almost complex structure. Likewise, $d_{C''}$ is defined as a Floer-theoretic differential on $\text{CF}(\widehat{V}, \widehat{V}')$. The maps c and h were constructed in Section 4.
- From Proposition 4.7, we have a canonical bijection between $\widehat{V} \cap \widehat{V}'$ and $\text{Fix}(\mu) \amalg \text{Fix}(\mu)$. We assign \mathbb{R} -degrees to elements of $\widehat{V} \cap \widehat{V}'$ so that the map (h_0, c_0) induced by this bijection preserves degree.
- We choose $\epsilon > 0$ as follows. First, as a consequence of Gromov–Floer compactness, there is some $\epsilon_0 > 0$ so that d_C , $d_{C''}$ and b all have order $[2\epsilon_0, \infty)$. There is ϵ_1 so that A'' has gap $(0, 2\epsilon_1)$. By Lemma 4.14, there is an $\epsilon_3 > 0$ so that, for any pair (u, v) contributing to c or h which has action $\leq 2\epsilon_3$, u is a gradient section and v is horizontal. Finally, there is ϵ_4 so that any gradient section v has action $A(v) \leq \epsilon_4/2$. We take $\epsilon < \min(\epsilon_0, \epsilon_1, \epsilon_2, \epsilon_3, \epsilon_4)$.

Now let us check the conditions of the double mapping cone lemma (5.4). The assumptions GAP and ORD 1 hold by choice of ϵ . The differentials $d_C = d_{C'}$ and $d_{C''}$ and the map b are positive by construction. Because $\epsilon < \epsilon_3$, Lemma 4.14 tells us that c_0 and h_0 precisely capture the low (meaning $< \epsilon$) action contributions to c and h , i.e. LOW 1 holds. It follows that h and c are positive, so POS holds.⁷ Moreover, ORD

⁷This explains why we chose to concentrate A and A' in degree 0 rather than grading them by representatives the action functional.

2 holds because $\epsilon < \epsilon_4$. The crucial condition LOW 2 comes from Lemmas 4.7, 4.11 and 4.12.

This completes the proof of Theorem 1. The first item in Addendum 1.6 is, at this point, a triviality: it is clear that the maps e and $\rho^!$ intertwine the action by quantum cap product by $QH^*(N)$ (in the case of e , this is just the fact that the cap product on fixed point Floer homology is an action of the small quantum cohomology algebra, and that this algebra is super-commutative). Hence so does the connecting map ρ_* . The remaining item of this addendum, concerning orientations, will be addressed in the next section.

As to Addendum 1.7, suppose that $\mu = \text{id}$ and that $\text{HF}(\widehat{V}, \widehat{V}) \cong H_*(\widehat{V}; \Lambda_R)$. We perturb μ to a Hamiltonian automorphism ϕ_H , generated by a C^2 -small Hamiltonian H . Then $\text{Fix}(\mu) = \text{crit}(H)$, and the Floer complex $\text{CF}(\phi_H)$ is canonically identified with the Morse complex for H . Similarly, $\text{CF}(\widehat{V}, (\phi_K \times \phi_H)\widehat{V})$ is identified with the Morse complex for $H \times K$ restricted to V . We have established that the low-action contributions to $C\rho^!$ and e are Morse-theoretic; they compute the transfer map $\rho^!$ on Morse homology, and the cap product by $e(V)$ on Morse homology. This gives the result.

6 Refinements

6.1 On the borderline

We now analyse the breakdown of the Gysin sequence in the borderline case where MON holds but MAS just fails because $m_{\widehat{V}}^{\min} = k + 1$. We saw at the outset (Example 1.10) that Theorem 1 can then also fail.

Consider the moduli space $\widetilde{\mathcal{N}}$ of parametrized J -holomorphic discs $\delta: (D, \partial D) \rightarrow (M_- \times N, \widehat{V})$ of Maslov index $k + 1$.

Because index $k + 1$ discs have the smallest positive action among holomorphic discs attached to \widehat{V} , $\mathcal{N} := \widetilde{\mathcal{N}}/\text{Aut}(D)$ is compact. For generic J , it is also smooth of dimension $(k + 1) + (2n - k) = 2n + 1$. According to Kwon–Oh [16], such discs must also be somewhere injective, which means that ‘transversality of evaluation’ holds: if we fix smooth, closed submanifolds A and B of \widehat{V} (or more generally smooth singular cycles), of respective codimensions α and β , the evaluation maps ev_1 and $\text{ev}_{-1}: \widetilde{\mathcal{N}} \rightarrow \widehat{V}$ are, for generic J , transverse to A and B respectively. Thus $\text{ev}_1^{-1}(A) \cap \text{ev}_{-1}^{-1}(B)$ is smooth of dimension $2n + 1 - \alpha - \beta$. Its quotient by the subgroup

$\mathbb{R} \subset \text{Aut}(D) = \text{PSU}(1, 1)$, consisting of those $\alpha \in \text{Aut}(D)$ with $\alpha(1) = 1$ and $\alpha(-1) = -1$, is compact. The action is free, since translation-invariant holomorphic discs are trivial, and so $\text{ev}_1^{-1}(A) \cap \text{ev}_{-1}^{-1}(B)/\mathbb{R}$ is compact, and smooth of dimension $2n - \alpha - \beta$.

In particular, if we take $A = \{x\}$ and $B = \rho^{-1}Z$, where Z represents the Poincaré dual to $e(V)$, then $\dim \text{ev}_1^{-1}(A) \cap \text{ev}_{-1}^{-1}(B)/\mathbb{R} = 2n - (2n - k) - (k + 1) = -1$, so the intersection is empty.

Now consider a global angular chain Y bounding $\rho^{-1}(Z)$ (as in Subsection 4.1.1). The space $(\text{ev}_1^{-1}(\{x\}) \cap \text{ev}_{-1}^{-1}(Y))/\mathbb{R}$ is a compact 0-manifold.

Definition 6.1 Define $\nu_Y = \#((\text{ev}_1^{-1}(x) \cap \text{ev}_{-1}^{-1}(Y))/\mathbb{R}) \in \mathbb{Z}/2$.

Note that ν_Y is independent of x and of (regular) J . For example, $\nu_Y = 0$ if, for some J , there is a point $x \in \widehat{V}$ which does *not* lie on a J -holomorphic disc of index $k + 1$. We do not address the dependence of ν_Y on Y .

Theorem 6.2 Suppose that MON holds and $m_{\widehat{V}}^{\min} = k + 1$. The conclusions of Theorem 1 hold when one replaces the map e by $e + t^{(k+1)\lambda} \nu_Y \text{id}$, where λ is the monotonicity constant of \widehat{V} .

Proof The construction of the map e obviously goes through. Less obviously, so does that of $\rho^!$: limits of sequences of sections of index 0 or 1 cannot bubble off discs, since these have index $\geq k + 1 \geq 2$. We will attempt to show that $\rho^! \circ e$ is nullhomotopic and see that we encounter an obstruction.

The first chain homotopy, H_1 , can still be defined as a map, because the relevant sections have index k . When we come to prove that H_1 is a homotopy, a potential problem arises: in the relevant moduli spaces \mathcal{Z} of index $k + 1$ sections, a disc $\delta: (D, \partial D) \rightarrow (M_- \times N, \widehat{V})$ of Maslov index $k + 1$ may bubble off. Indeed, \mathcal{Z} has extra ends parametrising limiting objects which consist of

- a pair (u, v) of pseudo-holomorphic sections of index 0; and
- a bubble $\delta: (D, \partial D) \rightarrow (M_- \times N, \widehat{V})$ such that $\delta(1) = v(p)$ and $\delta(-1) \in \text{im } \rho^{-1}Z$.

However, for generic almost complex structures,

$$\text{ev}_1^{-1}(v(p)) \cap \text{ev}_{-1}^{-1}(\rho^{-1}Z) = \emptyset,$$

so the problem is illusory and H_1 is a chain homotopy after all.

The second chain homotopy, H_2 , is constructed via index k sections (again, no bubbles). When try to prove that it is a homotopy by analysing a 1-dimensional moduli space we must again consider index $k + 1$ sections, so bubbling is possible. The moduli space now has ends corresponding to configurations consisting of:

- A pair (u, v) of pseudo-holomorphic sections of index 0; and
- a bubble $\delta: (D, \partial D) \rightarrow (M_- \times N, \widehat{V})$ such that $\delta(1) = v(p)$ and $\delta(-1) \in \text{im } Y$.

These configurations contribute the following boundary components to the 1-dimensional moduli space:

$$(61) \quad \bigcup_{(x_-, x_+) \in \mathcal{Z}(x_-, x_+)^0} \bigcup_{(u, v) \in \mathcal{Z}(x_-, x_+)^0} \mathcal{N}(v(p), Y).$$

This follows from the gluing theorem for holomorphic discs proved in Theorem [3, 4.1.2]. When $\nu_Y = 0$, there is for each $x_- \in \text{Fix}(\nu)$ and each $x_+ \in \widehat{V} \cap \widehat{V}'$, an even number of extra boundary components in the part of the moduli space consisting of sections asymptotic to (x_-, x_+) . Hence $h := H_1 + H_2$ is still a nullhomotopy. In general, the appearance of the boundary (61) implies that $\partial h + h\partial = \rho^! \circ e + \nu_Y t^{(k+1)\lambda} \rho^!$ (the Novikov weight $\lambda(k+1)$ is the area of the minimal Maslov-index discs). That is, h is a nullhomotopy of $\rho^! + t^{\lambda(k+1)} \nu_Y \text{id}$.

The analysis of low-action sections goes through unchanged. We may assume that $(k+1)\lambda \gg \epsilon$, which means that the new term $\nu_Y t^{(k+1)\lambda} \text{id}$ counts as a ‘high-action’ term. The algebraic part of the proof then goes over unchanged. \square

6.2 Compatibility with quantum cap product

The quantum cohomology $QH^*(N) = H^*(N; \Lambda_{\mathbb{Z}/2})$ acts on $\text{HF}_*(\mu)$ by quantum cap product. This makes it a module over the algebra $QH^*(N)$ with its quantum product. Since the map e is itself defined as the quantum cap product by the Euler class of $\rho: V \rightarrow N$, and since $QH^*(N)$ is (super) commutative, e intertwines the action of $QH^*(N)$.

The Lagrangian Floer homology $\text{HF}(\widehat{V}, \widehat{V}')$ is a module over $QH^*(M_- \times N)$, and hence also over $QH^*(N)$ via the pullback $\text{pr}_2^*: QH^*(N) \rightarrow QH^*(M_- \times N)$ (which is an algebra homomorphism). It follows easily from the definitions that $\rho^!$ intertwines the actions of $QH^*(N)$.

6.3 Orientations

To define Floer homology over $\Lambda_{\mathbb{Z}}$, one needs to specify a system of *coherent orientations*. This can be reduced to topology: to specify coherent orientations for the Lagrangian Floer homology $\mathrm{HF}_*(L_0, L_1)$ of $L_0, L_1 \subset P$, it suffices to give *relative spin-structures* on the two Lagrangians, that is, to give a stable vector bundle $\xi \rightarrow P$, together with stable spin-structures on $\xi|_{L_i} \oplus TL_i$ for $i \in \{0, 1\}$ [13, Chapter 6], [32].

The fixed point groups $\mathrm{HF}_*(\mu)$ are defined over $\Lambda_{\mathbb{Z}}$ in a canonical way, as explained in [24, II, Section 3.7].

In the setting of Lagrangian matching conditions, one wants the cobordism-maps to be well-defined over $\Lambda_{\mathbb{Z}}$. This turns out to be rather straightforward: one takes a vector bundle $\eta \rightarrow S_{\Gamma}$ over the surface S_{Γ} , and considers its restrictions $\eta|_{\Gamma_+}$ and $\eta|_{\Gamma_-}$ as bundles over Γ . Take $\xi \rightarrow \Gamma$ to be their direct sum $\xi = \eta|_{\Gamma_+} \oplus \eta|_{\Gamma_-}$. Giving a spin-structure for $TQ \oplus \xi|_Q$, where Q is the Lagrangian boundary condition. This choice defines a lift of the cobordism-maps over $\Lambda_{\mathbb{Z}}$ (we refer to [38] for a full discussion). Note that one can always work with the canonically-defined groups $\mathrm{HF}_*(\mu)$ (again, cf. [24, II, Section 3.7]).

When $V \subset M$ is an S^k -fibred coisotropic submanifold, the natural identification $\widehat{V} \rightarrow V$ (by projection to M) induces an isomorphism $T\widehat{V} \cong TV$ which we shall consider as an equality. If we put $\xi = p_1^*TM \rightarrow M \times N$, we will have $\xi|_{\widehat{V}} \oplus T\widehat{V} = TM|_V \oplus TV$. Thus there is a canonical short exact sequence

$$0 \rightarrow TV \oplus TV \rightarrow \xi|_{\widehat{V}} \oplus T\widehat{V} \rightarrow N_{V/M} \rightarrow 0$$

(the sequence splits as soon as one chooses a Riemannian metric on M and so identifies $N_{V/M}$ with TV^{\perp}). Now, a direct sum $U \oplus U$ of two copies of the same $O(r)$ -vector bundle always has a canonical spin-structure because its structure group reduces to the diagonal subgroup $O(r) \subset \mathrm{SO}(2r)$, and the inclusion homomorphism $O(r) \rightarrow \mathrm{SO}(2r)$ lifts to $\mathrm{Spin}(2r)$ since it kills π_1 . Applying this last principle to $U = TV$, we see that to give a spin-structure on $\xi|_{\widehat{V}} \oplus T\widehat{V}$ it suffices to give a spin-structure σ in $N_{V/M}$.

Now, σ also induces a spin-structure $(\mathrm{id} \times \mu)^*\sigma$ on the normal bundle to $\widehat{V}' = (\mathrm{id} \times \mu)\widehat{V}$, and hence gives rise to stable spin structures on both $\xi|_{\widehat{V}} \oplus T\widehat{V}$ and $\xi|_{\widehat{V}'} \oplus T\widehat{V}'$. Thus, given σ , the Floer homology $\mathrm{HF}(\widehat{V}, \widehat{V}')$ is canonically defined with Λ_R -coefficients for R an arbitrary commutative ring.

Lemma 6.3 *When V is S^1 -fibred, a co-orientation for V induces a relative spin structure for \widehat{V} . When V is the unit-sphere bundle $S(E)$ in a Euclidean vector bundle $E \rightarrow N$, a spin-structure in E induces a relative spin structure for \widehat{V} .*

Proof In the S^1 -fibred case, a co-orientation trivialises $N_{V/M}$, and so gives a canonical spin structure on $N_{V/M}$. When $V = S(E)$, we have $T^v V \oplus \mathbb{R} \cong \rho^* E$, and hence $w_2(N_{V/M}) = w_2(T^v V) = \rho^* w_2(E)$. By the classical Gysin sequence, this class vanishes if and only if $\text{rank}(E) = 2$ or $w_2(E) = 0$. In the latter case, a spin-structure in E induces a stable spin-structure on $T^v V$. But $N_{V/M} \cong (T^v V)^*$ via the symplectic form, so we get a stable spin-structure in $N_{V/M}$. \square

Our conclusion (which is the second item of Addendum 1.6) is as follows.

Proposition 6.4 *Suppose that the hypotheses of Theorem 1 hold. Assume either that V is a co-orientable hypersurface in M , or that $V = S(E)$ with $w_1(E) = 0 = w_2(E)$. Then $\text{HF}(\widehat{V}, (\text{id} \times \mu)\widehat{V})$ can be defined over Λ_R for any commutative unital ring R in such a way that the conclusions of Theorem 1 hold.*

(The bundle $\eta \rightarrow S_\Gamma$ required for the definitions of $C\rho$ and h can be specified by putting $\eta|_B = T^v Q$ and $\eta|_C = 0$.)

7 Connected sums of three-manifolds

This final section explains the (conjectural) connection between the symplectic Gysin sequence and gauge theory on 3- and 4-manifolds.

7.1 Connected sums and indefinite singularities

Let Y_1 and Y_2 be closed, oriented smooth 3-manifolds. When $f_1: Y_1 \rightarrow S^1$ and $f_2: Y_2 \rightarrow S^1$ are *harmonic Morse functions*, that is, circle-valued Morse functions with only indefinite critical points, the connected sum $Y_1 \# Y_2$ inherits a harmonic Morse function with $(\#\text{crit}(f_1) + \#\text{crit}(f_2) + 2)$ critical points. Indeed, let $I \subset S^1$ be a closed interval containing only regular points of f_1 and f_2 . Trivialise the two fibrations over I , identifying them with $I \times \Sigma_1 \rightarrow I$ and $I \times \Sigma_2 \rightarrow I$. There is an elementary cobordism W from $\Sigma_1 \amalg \Sigma_2$ to $\Sigma_1 \# \Sigma_2$, carrying a Morse function with a single critical point c_2 , of index 2. Likewise, there is an elementary cobordism W' from $\Sigma_1 \# \Sigma_2$ to $\Sigma_1 \amalg \Sigma_2$, carrying a Morse function with a single critical point c_1 of index 1. Glue W to W' along $\Sigma_1 \# \Sigma_2$ to obtain a 3-manifold V with a Morse function $V \rightarrow I$. Define a closed 3-manifold by gluing V into $f_1^{-1}(S^1 \setminus \text{int}(I)) \amalg f_2^{-1}(S^1 \setminus \text{int}(I))$ in the obvious way. Denote this new manifold by $Y_1 \#_I Y_2$. It comes with a Morse function

$f_1 \#_I f_2: Y_1 \#_I Y_2 \rightarrow S^1$, with two ‘new’ critical points, both over $\text{int } I$. It is easy to check that $Y_1 \#_I Y_2$ is diffeomorphic to the connected sum $Y_1 \# Y_2$; there is a disconnecting 2-sphere $S \subset Z$ lying over the interval I , the union of the ascending disc of c_2 and the descending disc of c_1 .

There is a natural cobordism X from $Y_1 \amalg Y_2$ to $Y_1 \#_{[a,b]} Y_2$, together with a function $F: X \rightarrow S^1 \times [0, 1]$ such that $F|_{(Y_1 \amalg Y_2)} = (f_1 \amalg f_2, 0)$ and $F|_{Y_1 \#_{[a,b]} Y_2} = (f_1 \#_{[a,b]} f_2, 1)$. To build it, begin with W . This carries a Morse function q with just one critical point. We normalise q so that $q(\Sigma_1 \amalg \Sigma_2) = \{0\}$, $q(\Sigma_1 \# \Sigma_2) = \{1\}$, and $q(W) = [0, 1]$. Let $H = \{z : |z| \leq 1, \text{Im}(z) \geq 0\} \subset \mathbb{C}$ be a closed half-disc, and let $m: H \rightarrow [0, 1]$ be the modulus function, $z \mapsto |z|$. Let $X_0 = m^*W$; so X_0 carries a natural map $F_0: X_0 \rightarrow H$.

Take the trivial cobordism $(Y_1 \amalg Y_2) \times [0, 1]$, and let $F_1 = (f_1 \amalg f_2, \text{id}): (Y_1 \amalg Y_2) \times [0, 1] \rightarrow S^1 \times [0, 1]$. Embed H into $[a, b] \times [0, 1]$ by an affine linear map so that its straight edge is mapped to $[a, b] \times \{1\}$. Thereby embed H into $S^1 \times [0, 1]$. Trivialise F_1 over H (extending the existing trivialisation over the straight edge). Now define X by excising $F_1^{-1}(H)$ from $(Y_1 \amalg Y_2) \times [0, 1]$ and gluing in X_0 in its place, in a way which should be clear. Thus

$$X = X_0 \cup ((Y_1 \amalg Y_2) \times [0, 1] \setminus F_1^{-1}(H)).$$

The map $F: X \rightarrow S^1 \times [0, 1]$ is obtained by gluing F_0 and (the restriction of) F_1 .

The function $\text{pr}_2 \circ F: X \rightarrow [0, 1]$ has just one critical point, namely, the unique point in the arc $\text{crit}(F_0)$ which lies over the image of $i \in H$ under its embedding into $S^1 \times [0, 1]$. This critical point c is non-degenerate of index 1. Hence X is the elementary cobordism from the disjoint union to the connected sum.

The critical manifold of the map $F: X \rightarrow S^1 \times [0, 1]$ is an ‘indefinite fold’, mapping injectively to the base; the local model is $(t; x_1, x_2, x_3) \mapsto (t; x_1^2 + x_2^2 - x_3^2)$. This makes it a simple example of a *broken fibration* in the sense of [24].

We shall concentrate on the case where f_1 and f_2 both have empty critical set, hence are fibrations. In this case, the set of critical values of F is an arc in the cylinder: it resembles the arc Γ from Figure 3.

7.2 Symmetric products

This section describes a class of examples of the symplectic Gysin sequence. In these examples, M is $\text{Sym}^n(\Sigma)$, the n th symmetric product of a closed Riemann surface Σ , for some $n \geq 1$. The symplectic form ω_λ on M is a Kähler form representing

one of the classes $\eta_\Sigma + \lambda\theta_\Sigma$ with $\lambda > 0$, where η_Σ is Poincaré dual to the class of $\{z\} + \text{Sym}^{n-1}(\Sigma)$ (for any $z \in \Sigma$) and θ_Σ is the pullback by the Abel–Jacobi map of the principal polarisation of the Picard torus $\text{Pic}^n(\Sigma)$. By a standard formula, one has

$$c_1(TM) = \left(n + \frac{\chi(\Sigma)}{2}\right) \eta_\Sigma - \theta_\Sigma.$$

Since θ_Σ integrates trivially on spheres, one finds that M is monotone (negatively monotone) provided that $n + \chi(\Sigma)/2 > 0$ (< 0).

There is a rich supply of automorphisms of $\text{Sym}^n(\Sigma)$ coming from self-diffeomorphisms of Σ . Indeed, if we fix an area-form α on Σ , the group of area-preserving diffeomorphisms acts on the symplectic mapping class groups of the symmetric products $\text{Sym}^n(\Sigma)$; there is a natural homomorphism

$$(62) \quad \kappa_{n,\lambda}: (\text{Aut} / \text{Ham})(\Sigma, \alpha) \rightarrow (\text{Aut} / \text{Ham})(\text{Sym}^n(\Sigma), \omega_\lambda).$$

This homomorphism was first constructed by gauge-theoretic means by D. Salamon [27], but it can also be described in elementary terms [24, II: Section 1].

When $\mu = \kappa_{\lambda,n}(\phi)$ for some area-preserving diffeomorphism $\phi \in \text{Aut}(\Sigma)$, the fixed-point Floer homology $\text{HF}(\mu)$ was conjectured by D. Salamon to be isomorphic to a version of Seiberg–Witten monopole Floer homology for $Y := T_\phi$ [27]. Salamon’s conjecture can be rephrased, using the calculations of [23], as follows.

Conjecture 7.1 *Take a Riemann surface Σ with a positive area form α , normalised to have area 1, inducing a closed 2-form α_ϕ on T_ϕ . Then there is an isomorphism of $\Lambda_{\mathbb{Z}/2}$ -modules*

$$(63) \quad \text{HF}_*(\kappa_{k,\lambda}(\phi)) \cong \bigoplus_{\langle c_1(\mathfrak{s}), [\Sigma] \rangle = 2(n+1-g)} \text{HM}_*(T_\phi, \mathfrak{s}; c_\lambda) \quad (\lambda > 0),$$

where on the right-hand side, a representative of

$$c_\lambda := -8\pi^2(\lambda^{-1} + n)[\alpha_\phi] - 32\pi^2 c_1(T^\vee T_\phi) \in H^2(T_\phi; \mathbb{R})$$

defines a perturbation of the Chern–Simons–Dirac functional. Moreover, there is an isomorphism (63) induced by a quasi-isomorphism of the underlying chain complexes, such that the quantum cap product by η_Σ goes over to the ‘ U -map’ [14, Chapter 1] on HM_* , up to chain homotopy.

The conventions here are those of Kronheimer–Mrowka [14] (note that the absence of decorations on the symbol HM reflects the absence of reducibles)

Returning to symplectic geometry of $M = \text{Sym}^n(\Sigma)$: it was shown in [24] that one can allocate to any embedded circle $\gamma \subset \Sigma$ an S^1 -fibred coisotropic hypersurface $V_\gamma \subset M$, up to Hamiltonian isotopy. Its reduced space is symplectomorphic to $N := \text{Sym}^{n-1}(\Sigma_\gamma)$, where Σ_γ is the surface obtained by surgery along γ . The symplectic structure $\bar{\omega}_\lambda$ is a Kähler form representing $\eta_{\Sigma_\gamma} + \lambda\theta_{\Sigma_\gamma}$. Thus there is an S^1 -bundle

$$\rho: V_\gamma \rightarrow N; \quad \rho^*\bar{\omega}_\lambda = \omega_\lambda|_{V_\gamma}.$$

The Euler class $e(V)$ is Poincaré dual to the difference of the divisors $p + \text{Sym}^{n-2}(\Sigma_\gamma)$ and $q + \text{Sym}^{n-2}(\Sigma_\gamma)$, where p and q lie in the two respective discs glued to $\Sigma \setminus \gamma$ to form Σ_γ .

Suppose that γ is a *separating* circle. Then

$$N = \coprod_{k=0}^{n-1} \text{Sym}^k(\Sigma_1) \times \text{Sym}^{n-1-k}(\Sigma_2),$$

where Σ_γ has connected components Σ_1 and Σ_2 . Write \widehat{V}^k for the component of \widehat{V}_γ that lies over $\text{Sym}^k(\Sigma_1) \times \text{Sym}^{n-1-k}(\Sigma_2)$. It was shown in [24, II, Section 4] that strong negativity for \widehat{V}^k in $\text{Sym}^n(\Sigma)_- \times \text{Sym}^k(\Sigma_1) \times \text{Sym}^{n-1-k}(\Sigma_2)$ can be arranged, for all $k \in \{0, \dots, n-1\}$, when $n \leq \frac{1}{2} \min(g(\Sigma_1), g(\Sigma_2))$. From Theorem 5 we obtain the following.

Proposition 7.2 *Assume γ separating, and let Σ_1 and Σ_2 be the connected components of Σ_γ . Let $\mu = \kappa_{k,\lambda}(\phi_1) \times \kappa_{n-1-k,\lambda}(\phi_2)$. When $2n \leq \min(g(\Sigma_1), g(\Sigma_2))$, there is a canonical isomorphism*

$$\text{HF}_*(\widehat{V}, (\text{id} \times \mu)\widehat{V}) \cong H \text{ cone}(\eta_{\Sigma_1} \otimes 1 - 1 \otimes \eta_{\Sigma_2})$$

where $\eta_{\Sigma_1} \otimes 1 - 1 \otimes \eta_{\Sigma_2}$ acts as a degree (-2) endomorphism of $\text{CF}(\mu)$.

Proposition 7.3 *Suppose that Conjecture 7.1 is true. Then there is an isomorphism*

$$(64) \quad \text{HF}(\widehat{V}, (\text{id} \times \mu)\widehat{V}) \cong \bigoplus_{c_1(t)=2(n+1-g)} \text{HM}_*(Y_1 \# Y_2, t; \tilde{c}_\lambda),$$

where \tilde{c}_λ is the class on $Y_1 \# Y_2$ induced by the c_λ -classes on the Y_i .

Proof We invoke an unpublished theorem of T. Mrowka and P. Ozsváth in Seiberg–Witten theory. (There are closely related results due to S. K. Donaldson and M. Furuta in Yang–Mills theory [7].) The theorem is that there is a natural isomorphism

$$\text{HM}_*(Y_1 \# Y_2) \cong \text{cone}(U \otimes 1 - 1 \otimes U)$$

where $U \otimes 1 - 1 \otimes U$ acts as a degree (-2) endomorphism of $HM_*(Y_1) \otimes_{\Lambda_{\mathbb{Z}/2}} HM_*(Y_2)$. (We work over $\Lambda_{\mathbb{Z}/2}$ for convenient comparison with symplectic Floer homology.) The cobordism X induces a homomorphism $HM_*(X): HM_*(Y_1) \otimes_{\Lambda_{\mathbb{Z}/2}} HM_*(Y_2) \rightarrow HM_*(Y_1 \# Y_2)$, and the composite $HM_*(X) \circ (U \otimes 1 - 1 \otimes U)$ is nullhomotopic because the basepoints in Y_1 and Y_2 , used in the definitions of the U -maps, are homotopic in X . Any version of monopole Floer homology can be used, and in particular closed 2-form perturbations are allowed, provided that they extend over X .

In our case, we can take a near-symplectic form on X whose zero set is $\text{crit}(F)$ (cf. [24]). Comparing the Mrowka–Ozsváth isomorphism with Proposition 7.2 establishes the result. \square

The last proposition is closely related to the conjecture from [24] that the ‘Lagrangian matching invariants’ for broken fibrations defined there are equal to the Seiberg–Witten invariants of the underlying 4-manifolds.

7.3 Non-separating circles

A simpler situation occurs when $\gamma \subset \Sigma$ is a non-separating circle. Since $n + \chi(\Sigma)/2 = (n - 1) + \chi(\Sigma_\gamma)/2$, N is (anti-)monotone precisely when M is. The product $M_- \times N$ is also (anti-)monotone, with monotonicity constant $n + \chi(\Sigma)/2$, when this constant is positive (negative), and its minimal Chern number is

$$c_1^{\min}(M_- \times N) = |n + \frac{\chi(\Sigma)}{2}| = |n - 1 + \frac{\chi(\Sigma_\gamma)}{2}|.$$

The minimal Maslov index is [24]

$$m_{\widehat{V}}^{\min} = 2c_1^{\min}(M_- \times N) = |2n + \chi(\Sigma)| = |2(n - 1) + \chi(\Sigma_\gamma)|.$$

The circle-bundle $\rho: V_\gamma \rightarrow N$ has $e(V) = 0$, though there is no preferred trivialisation. Thus we can make the following deduction from Theorem 1 and its strongly negative counterpart.

Proposition 7.4 *Assume γ non-separating. Take any $\mu \in \text{Aut}(N, \bar{\omega}_\lambda)$, for example, $\mu = \kappa_{n-1, \lambda}(\phi)$. Fix a trivialisation of ρ . When $n > g(\Sigma)$ or $2n \leq g(\Sigma) - 1$, there is then a canonical isomorphism of relatively graded modules,*

$$\text{HF}(\widehat{V}, (\text{id} \times \mu)\widehat{V}) \cong \text{HF}(\mu) \oplus \text{HF}(\mu)[1].$$

This sequence also has a topological counterpart. Here one starts with the connected 3-manifold $Y = T_\phi$, and forms its *internal* connected sum by excising a pair of disjointly embedded discs and gluing in a 1-handle $S^2 \times [0, 1]$ to obtain a new 3-manifold Y' which has the diffeomorphism type of $Y \# (S^1 \times S^2)$. When Y is fibred over S^1 , by a map f , Y' maps to S^1 by a Morse function f with precisely two critical points, of indices 1 and 2, just as before. Moreover, there is an elementary cobordism from X from Y to Y' , and a map F (broken fibration) from X to an annulus, interpolating between f and f' , with an arc of critical points forming a non-degenerate critical manifold. The author does not know an explicit source in the literature for the behaviour of monopole Floer homology under 1-handle attachment, but the behaviour of another conjecturally isomorphic theory, Heegaard Floer homology is well-understood, and tallies with our result.

7.4 Flat connections

Just as symplectic geometry of symmetric products mimics Seiberg–Witten theory, so symplectic geometry of moduli spaces of flat connections models instanton Floer homology (cf. Dostoglou–Salamon [9]), at least in cases where reducible flat connections can be avoided.

Consider the twisted character variety M_g of a punctured genus $g \geq 1$ surface:

$$M_g = \{(A_1, B_1; \dots; A_g, B_g) \in \mathrm{SU}(2)^{2g} : \prod_{i=1}^g [A_i, B_i] = -I\} / \mathrm{SU}(2).$$

Here $\mathrm{SU}(2)$ acts by conjugation. Among many possible references for the geometry of M_g , we mention [35], because it discusses the locus V to be introduced momentarily.⁸ The space M_g is a canonically smooth and symplectic manifold of dimension $6g - 6$. Assuming $g \geq 2$, the subspace $V := \{A_g = I\} \subset M_g$ has codimension 3, and there is a natural projection $\rho: V \rightarrow M_{g-1}$. This map is an S^3 -bundle with vanishing Euler class (the locus $\{B_g = A_g = I\} \subset M_g$ defines a section S). It is easily verified to be the reduction map of an S^3 -fibred coisotropic submanifold (to see this one can use either the Atiyah–Bott symplectic structure, thinking of the A_i and B_i as holonomies of flat connections, or Goldman’s finite-dimensional reformulation of it).

Because M_g and M_{g-1} are simply connected, one has

$$m_V^{\min} = 2c_1^{\min}(M_{g-} \times M_{g-1}) = 4 = 3 + 1,$$

⁸There are also interesting unpublished studies by M. Callahan [6] and P. Seidel [28] of the role of V in symplectic monodromy problems.

which puts us on the borderline for applicability of the Gysin sequence. The global angular chain (actually a cycle if we take the zero-cycle to represent the Euler class) may be taken to be the section S .

Regarding elements of M_g as representations of $\pi_1(\Sigma)$, where Σ is a genus g surface equipped with a standard basis of curves, one sees that there is a natural homomorphism $\theta: \pi_0 \text{Diff}^+(\Sigma) \rightarrow \text{Aut}(M_g, \omega)$, $\phi \mapsto (\phi^{-1})^*$. According to Theorem 6.2, one has

$$\text{HF}_*(\widehat{V}, \theta(\phi)\widehat{V}) \cong \text{cone}(\nu_S t^a \text{id})$$

for some $a > 0$. Thus $\text{HF}_*(\widehat{V}, \theta(\phi)\widehat{V})$ is either always (i.e., for all ϕ) zero, or else it is always isomorphic, as a relatively graded module, to $\text{HF}_*(\theta(\phi)) \oplus \text{HF}_*(\theta(\phi))[3]$.

However, $\text{HF}_*(\widehat{V}, \widehat{V})$ is non-zero for any $g \geq 2$: by a theorem of Albers [1, Corollary 2.11], a monotone Lagrangian $L \subset M^{2n}$ of minimal Maslov index ≥ 2 can only have vanishing Floer self-homology if $[L] = 0 \in H_n(M; \mathbb{Z}/2)$; but $[V]$ is non-zero in $H_*(M_g; \mathbb{Z}/2)$, hence $[\widehat{V}] \neq 0 \in H_*(M_g \times M_{g-1})$. We make the following conclusion.

Theorem 7.5 *A choice of section of $\rho: V \rightarrow M_{g-1}$ determines isomorphisms*

$$\text{HF}_*(\widehat{V}, \widehat{V}) \cong \text{HF}_*(\theta(\phi)) \oplus \text{HF}_*(\theta(\phi))[3]$$

for all $\phi \in \text{Diff}^+(\Sigma)$.

This appears to tally with the behaviour of $\text{SO}(3)$ -instanton Floer homology under internal connected sum, though we do not attempt to make the connection precise. It would be interesting to find a symplectic counterpart for the behaviour of external connected sums in instanton theory, as described in [12, 7].

References

- [1] P. Albers, *A Lagrangian Piunikhin–Salamon–Schwarz morphism and two comparison homomorphisms in Floer homology*, Int. Math. Res. Not. (2007), doi:10.1093/imrn/rnm134.
- [2] M. F. Atiyah, K. Patodi and I. Singer, *Spectral asymmetry and Riemannian geometry. I*, Math. Proc. Cambridge Philos. Soc. 77 (1975), 43–69
- [3] P. Biran and O. Cornea, *Quantum structures for Lagrangian submanifolds*, 2007 preprint, ArXiv: 0708.4221v1 [math.SG].
- [4] P. Biran and M. Khanevsky, *A Floer–Gysin exact sequence for Lagrangian submanifolds*, in preparation.

- [5] R. Bott and L. Tu, *Differential forms in algebraic topology*, Graduate Texts in Mathematics, 82. Springer-Verlag, New York-Berlin, 1982.
- [6] M. Callahan, incomplete draft for D. Phil. thesis, Oxford, circa 1996.
- [7] S. K. Donaldson, *Floer homology groups in Yang–Mills theory*, Cambridge Tracts in Mathematics, 147. Cambridge University Press, Cambridge, 2002.
- [8] S. K. Donaldson and P. B. Kronheimer, *The geometry of four-manifolds*, Oxford Mathematical Monographs. Oxford Science Publications. The Clarendon Press, Oxford University Press, New York, 1990.
- [9] S. Dostoglou and D. Salamon, *Self-dual instantons and holomorphic curves*, Ann. of Math. (2) 139 (1994), no. 3, 581–640.
- [10] A. Floer, *The unregularized gradient flow of the symplectic action*, Comm. Pure Appl. Math. 41 (1988), no. 6, 775–813.
- [11] A. Floer, H. Hofer and D. A. Salamon, *Transversality in elliptic Morse theory for the symplectic action*, Duke Math. J. 80 (1995), no. 1, 251–292.
- [12] K. Fukaya, *Floer homology of connected sum of homology 3-spheres*, Topology 35 (1996), no. 1, 89–136.
- [13] K. Fukaya, Y.-G. Oh, H. Ohta and K. Ono, *Lagrangian intersection Floer theory: anomaly and obstruction*, book manuscript, 2000.
- [14] P. B. Kronheimer and T. S. Mrowka, *Monopoles and three-manifolds*, New Mathematical Monographs, 10, Cambridge University Press, Cambridge, 2007.
- [15] H. Hofer and D. A. Salamon, *Floer homology and Novikov rings*, The Floer memorial volume, 483–524, Progr. Math., 133, Birkhäuser, Basel, 1995.
- [16] D. Kwon and Y.-G. Oh, with an appendix by J.-P. Rosay, *Structure of the image of (pseudo)-holomorphic discs with totally real boundary condition*, Comm. Anal. Geom. 8 (2000), no. 1, 31–82.
- [17] L. Lazzarini, *Existence of a somewhere injective pseudo-holomorphic disc*, Geom. Funct. Anal. 10 (2000), no. 4, 829–862.
- [18] D. McDuff and D. Salamon, *Introduction to symplectic topology*, 2nd edition, Oxford Mathematical Monographs, Oxford University Press, New York, 1998.
- [19] D. McDuff and D. Salamon, *J-holomorphic curves in symplectic topology*, American Mathematical Society Colloquium Publications, 52. American Mathematical Society, Providence, RI, 2004.
- [20] J. Milnor, *Morse theory*, Annals of Mathematics Studies, No. 51. Princeton University Press, Princeton, N.J., 1963.
- [21] Y.-G. Oh, *Floer cohomology of Lagrangian intersections and pseudo-holomorphic disks, I*, Comm. Pure Appl. Math. 46 (1993), no. 7, 949–993; Addendum, Comm. Pure Appl. Math. 48 (1995), no. 11, 1299–1302.

- [22] Y.-G. Oh, *Floer cohomology, spectral sequences, and the Maslov class of Lagrangian embeddings*, Internat. Math. Res. Notices 1996, no. 7, 305–346.
- [23] T. Perutz, *Symplectic fibrations and the abelian vortex equations*, Commun. Math Phys. 278, 289–306 (2008).
- [24] T. Perutz, *Lagrangian matching invariants for fibred four-manifolds: I*, Geom. Topol. 11 (2007), 759–828; *II*, Geom. Topol. 12 (2008), 1461–1542.
- [25] W.-D. Ruan, *Deformation of integral coisotropic submanifolds in symplectic manifolds*, J. Symplectic Geom. 3 (2005), no. 2, 161–169.
- [26] D. Salamon, *Lectures on Floer homology*, Symplectic geometry and topology (Park City, UT, 1997), 143–229, IAS/Park City Math. Ser., 7, Amer. Math. Soc., Providence, RI, 1999.
- [27] D. Salamon, *Seiberg-Witten invariants of mapping tori, symplectic fixed points, and Lefschetz numbers*, Proceedings of 6th Gökova Geometry–Topology Conference. Turkish J. Math. 23 (1999), no. 1, 117–143.
- [28] P. Seidel, *Symplectic monodromy*, unpublished notes, 1997.
- [29] P. Seidel, *More about vanishing cycles and mutation*, Symplectic geometry and Mirror Symmetry (Seoul, 2000), 429–465, World Sci. Publ., River Edge, NJ, 2001.
- [30] P. Seidel, *A long exact sequence for symplectic Floer cohomology*, Topology 42 (2003), no. 5, 1003–1063.
- [31] P. Seidel, *Lectures on 4-dimensional Dehn twists*, in *Symplectic 4-Manifolds and Algebraic Surfaces. Lectures given at the C.I.M.E. Summer School held in Cetraro, Italy, September 2-10, 2003*, Lecture Notes in Mathematics, 1938. Springer, Berlin, 2008.
- [32] P. Seidel, *Fukaya categories and Picard–Lefschetz theory*, book manuscript, 2006. To appear in ETH Lecture Notes Series (European Math. Soc.).
- [33] J.-P. Serre, *Homologie singulière des espaces fibrés*, Ann. of Math. (2) 54, (1951), 425–505.
- [34] S. Smale, *On the structure of 5-manifolds*, Ann. of Math. (2) 75 (1962), 38–46.
- [35] M. Thaddeus, *A perfect Morse function on the moduli space of flat connections*, Topology 39 (2000), no. 4, 773–787.
- [36] C. Weibel, *An introduction to homological algebra*, Cambridge Studies in Advanced Mathematics, 38. Cambridge University Press, Cambridge, 1994.
- [37] K. Wehrheim and C. Woodward, *Functoriality for Lagrangian correspondences in Floer theory*, 2007 preprint, arXiv: 0708.2851 [math.SG].
- [38] K. Wehrheim and C. Woodward, *Orientations for pseudo-holomorphic quilts*, in preparation.